

Optimizing Dell EMC SC Series Storage for Oracle OLAP Processing

Abstract

This paper highlights the advanced features of Dell EMC™ SC Series storage and provides guidance on how they can be leveraged to deliver a cost-effective solution for Oracle® OLAP and DSS deployments.

November 2019

Revisions

Date	Description
September 2014	Initial release
October 2016	Updated for agnosticism with Dell SC Series all-flash arrays; updated format
July 2017	Consolidated best practice information and focus on OLAP environment; updated format
November 2019	vVols branding update

Acknowledgements

Updated by: Mark Tomczik and Henry Wong

The information in this publication is provided “as is.” Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

© 2014–2019 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

Dell believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Table of contents

1	Introduction.....	5
1.1	Audience.....	5
1.2	Prerequisites.....	5
2	SC Series overview.....	6
3	Optimize an SC Series array for an Oracle data warehouse.....	8
3.1	Create a high availability SAN.....	8
3.1.1	Redundant hardware components, SC Series virtual ports, and fault domains.....	8
3.1.2	Multipath I/O software.....	8
3.2	Direct attach with front-end SAS.....	9
3.3	Storage layout for an Oracle database.....	9
3.3.1	Oracle ASM.....	11
3.3.2	Linux LVM.....	12
3.4	Achieve greater storage efficiency and performance.....	12
3.4.1	Storage tiers and tier redundancy.....	12
3.4.2	Dynamic data placement with storage profiles.....	13
3.4.3	Initial loading of historical data.....	16
3.4.4	Data reduction.....	16
3.4.5	Thin provisioning compared to preallocating storage.....	19
3.4.6	Volume cache.....	20
3.5	Protecting Oracle with snapshots and a consistency group.....	20
3.6	Sizing considerations.....	21
3.6.1	Collect and measure performance statistics.....	21
3.6.2	Configure I/O for bandwidth.....	22
3.6.3	Capacity planning.....	22
3.6.4	Test the I/O system before implementation.....	23
4	Conclusion.....	24
A	Storage profiles and RAID levels.....	25
B	Technical support and resources.....	27

Executive summary

Intelligent choices must be made when selecting the hardware for an environment that will service the needs of a business and expected data growth. When talking about data warehouses, higher host processing capabilities coupled with high-performing storage solutions such as Dell EMC™ SC Series arrays deliver the power and throughput necessary to meet the demands of business intelligence and reporting applications that provide trend analysis, forecasting, and business analytics.

Typical data warehouse systems move large amounts of sequential data in large multi-block read/write I/O streams and tend to be very read intensive. The I/O streams can also be 100% read-only during periods of report generation, but then change to an I/O mix during the extract, transform, and load (ETL) stage. SC Series arrays are an excellent choice for data warehousing because of the Dell™ Fluid Data™ architecture that keeps the data as dynamic as the organization it supports by placing the right data in the right place at the right time, regardless of the I/O mix of the application.

SC Series arrays help resolve many of the storage challenges facing IT organizations by providing many advanced features typically only found in expensive high-end storage. This paper highlights some of these features and provides guidance on how they can be leveraged to deliver a cost-effective solution for the demanding performance requirements of Oracle OLAP and DSS deployments.

1 Introduction

The purpose of this document is to provide general best practices and sizing guidelines when deploying an Oracle data warehouse on any SC Series array. The information provided in this document is built upon the guiding principles and building blocks in the document, [Dell EMC SC Series Arrays and Oracle](#). Therefore, it should be thoroughly reviewed before implementing the best practices mentioned in this document. Actual configuration and sizing of the components will vary based on individual business, application, and infrastructure requirements. The recommendations presented in this document may not be relevant to all environments and should be evaluated against specific business needs and adjusted accordingly.

This document is not intended to provide a step-by-step configuration or be an exact sizing guide. It also should not be considered an exhaustive or authoritative source on any component discussed. Performance tuning of any component within the document is also not covered. For information on SC Series arrays, see the *Dell Storage Manager Administrator's Guide* available on the [Dell Support site](#). For detailed information on any part of Oracle configuration, use the information provided at the [My Oracle Support site](#). For detailed information on best practices on SC Series storage for Oracle deployments, see the resources listed in appendix B.

1.1 Audience

This document is intended for information technology professionals seeking to deploy a cost-effective Oracle OLAP and DSS database environment using an SC Series array.

1.2 Prerequisites

It is assumed the reader has some knowledge of the benefits of SC Series storage and basic sizing principles related to I/O subsystems, and understands key I/O performance metrics: IOPS, throughput, latencies, as well as random I/O versus sequential I/O. Information pertaining to those topics are covered in the document, [Dell EMC SC Series Arrays and Oracle](#), and should be reviewed before reading this document.

Readers should have prior experience with or training in:

- Data warehousing concepts
- Basic server technologies
- Basic operating system technologies
- Basic sizing principles related to I/O subsystems
- Understanding key I/O performance metrics: IOPS, throughput, latencies, as well as random I/O versus sequential I/O
- Validating I/O paths and storage designs
- Different RAID levels, Fibre Channel, iSCSI, serial-attached SCSI (SAS), and IP networking administration
- Multipath software
- Operation, configuration, and benefits of SC Series storage
- Using and administrating Oracle Linux, Oracle RDBMS 11g or above, and Oracle Automated Storage Management (ASM)
- Oracle standalone and RAC environments

2 SC Series overview

The SC Series array is built on the core Fluid Data architecture that has made flash storage affordable to customers. Traditional storage arrays typically can accommodate one storage type, or when multiple storage types are used, the storage must be configured in such a way that it is inflexible and difficult to adjust to any future workload changes. With the advanced storage virtualization technology built into the core of the Dell Storage Center OS (SCOS), SC Series arrays can fully utilize both flash and HDD media and automatically optimize and adapt to the changing workload on demand.

At the center of SC Series is SCOS, which incorporates the best of flash technology, powers the storage virtualization, and delivers greater flexibility and economics for both all-flash and hybrid arrays.

The following list highlights the core SCOS features. These powerful features are transparent to the OS and applications, such as Linux® and Oracle database.

Dynamic block architecture records and tracks metadata for every block of data and provides system intelligence on how those blocks are being used. The metadata enables SC Series storage to take a more sophisticated and intelligent approach to storing, recovering, and managing data.

Storage virtualization occurs at the disk level, creating a dynamic pool of storage resources shared by all servers. Because read/write operations are spread across all available drives within the same tier, multiple requests are processed in parallel, boosting system performance.

Dynamic capacity (thin provisioning) delivers high storage utilization by eliminating allocated but unused capacity. It completely separates storage allocation from utilization, enabling users to create any size of virtual volume upfront, and only consume actual physical capacity when data is written.

Snapshots (replays) provide continuous, space-efficient data protection. A snapshot taken of a volume creates a point-in-time copy (PITC) of the volume by making all written pages read-only. Any further changes to the volume get written to new pages (active data). When the volume is read, SC Series storage seamlessly presents the read-only pages from the snapshot and any active data. Consistent snapshots in an Oracle environment can be effective for database backups, recoveries, and cloning. For more information on consistent snapshots, see section 3.5.

Data Progression (DP) is a Fluid Data storage or automated tiered storage feature that automatically migrates data to the optimal storage tier based on a set of predefined or custom policies called storage profiles. Data Progression eliminates the need to manually classify and migrate data to different storage tiers while reducing the number and cost of drives and reducing the cooling and power costs.

Fast Track technology enhances automated tiered storage by dynamically placing the most frequently accessed data on the fastest, or outer, tracks of each hard disk drive. Fast Track does not require any manual configuration and it is licensed separately.

Deduplication and compression enable data reduction and work with Data Progression to reduce the data stored on the array. A Data Progression license includes these data-reduction features without extra cost. They are enabled on a per-volume basis. Data is processed for deduplication and compression during the scheduled daily Data Progression cycles or when snapshots are taken, which trigger an on-demand Data Progression on the volume. The feature is supported on all SC Series models except the SCv Series. For more information on SC Series data reduction, refer to the [Dell Storage Center OS 7.0 Data Reduction with Deduplication and Compression Solution Guide](#).

Replication technology protects a virtual volume by creating a replica on a remote SC Series array using either asynchronous or synchronous mode. After the initial setup, SCOS copies only the data changes made since the last snapshot. This uses less network bandwidth and less time. In the event of primary array/site failure, the replicas can be activated by the administrator on the remote site.

Live Volume is a hyper-volume that stretches across two different SC Series arrays. The feature extends data protection beyond the basic array replication and offers enhanced disaster recovery capability to a remote site. When the primary SC Series array becomes unavailable, Live Volume can seamlessly fail over to the secondary array without causing data access interruption. The feature allows on-demand failover or auto-failover between the two arrays. For more information on Live Volume and how it interacts with Microsoft® and VMware®, refer to [SC Series Technical Documents](#).

Live Migrate builds on the same virtualization architecture of Live Volume and allows migrating workloads seamlessly from one array to another. The feature is intended for migrating data between SC Series arrays without causing service interruption. It simplifies the complex manual migration workflow of setting up the Live Volume and replication, mapping the Live Volume to the server, waiting for replication to finish, performing role swap, and removing the Live Volume at the end.

Dell Storage Manager (DSM) serves as a centralized administration console for one or multiple SC Series arrays, PS Series arrays, and FluidFS clusters. DSM offers an easy-to-use graphical interface to manage storage volumes, server mappings, storage-medium configuration, storage profiles, snapshot profiles, replication, and Live Volumes, and provides performance monitoring of the storage systems with rich charting capabilities. DSM also provides product integration for Microsoft and VMware products. These products include Microsoft System Center Virtual Machine Manager, VMware Site Recovery Manager, and VMware® vSphere® Virtual Volumes™ (vVols). For more information, refer to the *Dell Storage Manager Administrator's Guide* and visit [SC Series Technical Documents](#).

SC Series PowerShell SDK, REST API, and Java-based CompCU allow administrators to manage the SC Series storage through scripting and automation. They are available to customers free of charge.

Dell SupportAssist is built into DSM and every SC Series array. When enabled, DSM sends array configuration, diagnostic, and performance data to Dell technical support for proactive monitoring, performance analysis, and troubleshooting purposes. With SupportAssist, administrators can choose to receive update notification and whether the updates should be pushed directly to DSM automatically. Administrator approval is always required to apply the updates.

For additional documentation on these features, see the *Dell Storage Manager Administrator's Guide* and the best practices document, [Dell EMC SC Series Arrays and Oracle](#).

3 Optimize an SC Series array for an Oracle data warehouse

Unlike an OLTP database, a data warehouse is characterized by relatively low volume of transactions. The majority of the activities in a data warehouse involves complex queries and aggregates large set of data. The volume of data tends to grow steadily over time and is kept available for a much longer period of time.

The following sections discuss ways to optimize an SC Series array and Oracle storage for a data warehouse environment.

3.1 Create a high availability SAN

A well-designed SAN should always include redundant hardware components to eliminate a single point of failure. SC Series arrays are designed with high availability (HA) in mind. Each array consists of two controllers and multiple front-end and back-end ports. SC Series arrays support simultaneous transport protocols including Fibre Channel (FC) and iSCSI. For improved performance in a data warehouse, it is recommended to use high speed FC connectivity and high speed (12Gb) back-end SAS connectivity.

3.1.1 Redundant hardware components, SC Series virtual ports, and fault domains

For each connection type, ensure redundant adapter cards are installed in each SC Series controller. To provide maximum HA protection and bandwidth optimization, configure the front-end ports on the SC Series array in **virtual port** mode and groups into multiple **fault domains** which provide port-level and controller-level failover functionality. Similarly, redundant SAN switches are preferred.

Figure 1 illustrates an example of FC connectivity in a dual fault domain configuration on an SC5020 array. Each domain consists of two FC ports from each controller. A similar setup can also apply to a dual iSCSI fault domain configuration. FC and iSCSI ports cannot be mixed in the same fault domain.

On an Oracle database server, configure a minimal of two FC or iSCSI ports for redundancy. To handle the large throughput and bandwidth required during the ETL process or the data ingestion, use multiple high-speed FC or iSCSI HBAs or ports such as 32Gb FC or 40Gb iSCSI.

It is not recommended to present an SC Series LUN to the same server on both the FC and iSCSI SAN because of different performance characteristics.

3.1.2 Multipath I/O software

Regardless of FC or iSCSI SAN use, multipath software is strongly recommended on the database server to provide I/O load balancing, bandwidth aggregation, and automatic path failover. SC Series arrays work seamlessly with various multipath software including Linux DM-Multipath, Dell EMC PowerPath™ and Microsoft Windows® MPIO.

If the database server runs on a virtual machine (VM) and the SC Series volumes are virtualized through a hypervisor such as VMware vSphere®, multipath software should be configured on the physical nodes and is not required in the database VM.




Figure 1 Redundant ports and paths, and HA fault domains

3.2 Direct attach with front-end SAS

For SC Series models that support using front-end SAS connectivity to the servers, configure multiple SAS connections between each server and the array to provide path redundancy. Multipath software should be configured on each server to provide automatic path failover and I/O load balancing. Sharing an SC Series volume through front-end SAS on multiple servers for an Oracle RAC configuration is not supported. For shared storage configuration in an Oracle RAC deployment, use an FC or iSCSI SAN to provide shared storage between the RAC cluster nodes.

3.3 Storage layout for an Oracle database

Figure 2 depicts an Oracle RAC database model configured on a SC Series array and managed by Oracle ASM. While the configuration will vary based on each environment's unique requirements, the guiding principles remain similar.

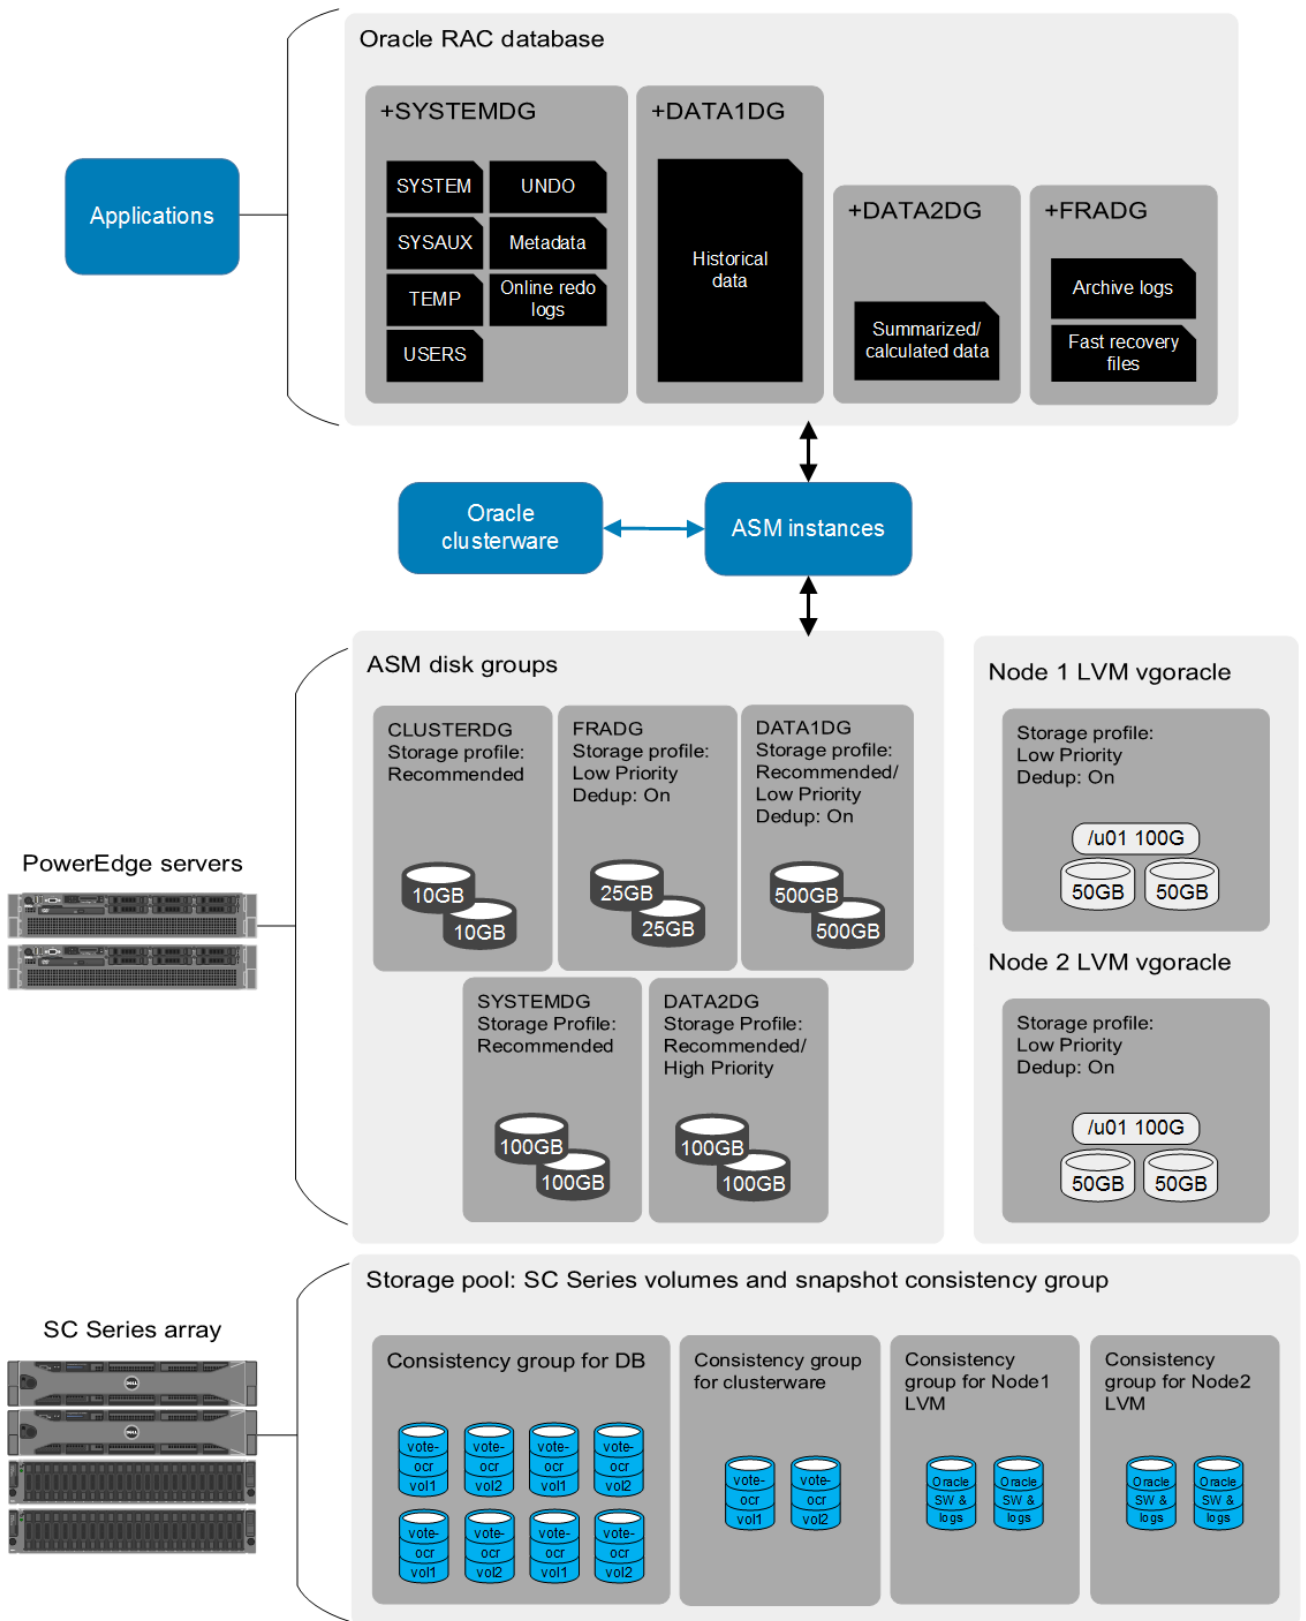


Figure 2 Oracle database storage layout model

3.3.1 Oracle ASM

Dell EMC and Oracle recommend using Oracle Automatic Storage Management (ASM) to manage SC Series volumes for the database and clusterware. This section reviews the general guidelines and additional considerations for a data warehouse.

ASM guidelines:

- Configure a separate diskgroup for the Oracle clusterware.
- A data warehouse can span across multiple diskgroups. Arrange different data types into separate disk groups where appropriate storage profiles and data reduction features can be applied to maximize storage efficiency and performance. See section 3.4 for further details.
- Make sure volumes in the same diskgroup have the same capacity and volume attributes such as data reduction, volume cache and storage profile on a SC Series array.
- Use fewer but larger volumes to reduce the number of objects to be managed. SCOS supports up to 2,000 volumes in a single SC Series array and a maximum volume size of 500TB.
- Make sure each disk group contains even numbers of volumes to allow I/O distribution to both SC Series controllers, hence, maximizing the performance and I/O bandwidth for the data warehouse.
- Ensure that diskgroups are mounted and used by a data warehouse exclusively and not shared with other databases. This simplifies using array-based snapshots for backup and recovery.
- In order to take an array-based snapshot on a multi-volume Oracle data warehouse, make sure that all volumes belonging to the same data warehouse are snapped together. This can be achieved by creating a consistent snapshot profile and assigning it to the volumes. In Figure 2, for example, volumes in +SYSTEMDG, +DATA1DG, +DATA2DG, and +FRADG have the same consistent snapshot profile.
- Keep the total number of ASM disks for a data warehouse (+SYSTEMDG, +DATA1DG, +DATA2DG and +FRADG) below the snapshot consistency group limits. The limits vary across SC Series array models and SCOS versions. For example, an SC5020 array with SCOS 7.2 supports 50 volumes per consistency group, while an SC7020 with SCOS 7.2 supports 100 volumes per consistency group. The limits are documented in the SCOS release notes (found on Dell.com/support) under the **Scalability Guidelines**. See section 3.5 for more details on snapshots and consistency groups.
- While ASM can provide software-level mirroring, it is not necessary because data protection has already been provided by the built-in SC Series multi-RAID levels. Use **External Redundancy** for ASM disk groups to enable substantial storage savings on a large data warehouse. This also reduces overall IOPS from ASM and results in better I/O performance.

For more information on Oracle ASM best practices with SC Series arrays, see the document, [Dell SC Series Storage with Oracle ASM Best Practices](#).

3.3.2 Linux LVM

Linux Logical Volume Manager (LVM) is a common general-purpose storage manager included in all popular Linux distributions. Since Oracle software is excluded from ASM, LVM can be used to store the software on a local Linux file system. Similar to ASM, it is not necessary to mirror logical volumes in LVM. LVM allows striping data files across multiple volumes in a volume group. Unlike ASM, the striping needs to be configured by administrators explicitly, and automatic restriping is not available when extending a volume group in LVM.

LVM guidelines:

- Create a dedicated LVM volume group for storing the Oracle software.
- Use two or more volumes (LUNs) in a volume group to increase performance.
- Make sure all volumes in the same volume group have the same volume attributes (such as data reduction, volume cache, and storage profile on the SC Series array)
- When multiple copies of Oracle software are installed on the server, create one volume group for each copy of Oracle software. This allows array-based snapshots on individual Oracle software copy.
- Set up a snapshot consistency profile for each multi-LUN volume group. Assign all volumes in the same volume group to the same snapshot consistency profile.
- In an Oracle RAC configuration, use a dedicated local LVM volume group for each cluster node.
- Enable the **discard** filesystem option in `/etc/fstab`, or include **-o discard** to the **mount** command. Certain filesystem types, such as ext4, supports the online TRIM command. This is useful for thin-provisioned volumes because when deleting files in the filesystem, the associated blocks are also deleted or released in the SC Series array.

Note: If an SC Series volume is a preallocated volume, to keep the preallocated space, do not enable the **discard** filesystem option or run the TRIM command against the file system. Otherwise the space will be deallocated.

3.4 Achieve greater storage efficiency and performance

A typical data warehouse consumes large amount of storage space to store many years of historical data. While flash drives offer a significant performance improvement and the cost has come down rapidly in recent years, they can be more expensive than HDDs, which can also offer a larger capacity per drive. One solution is to combine both flash drives and HDDs to address the need of a data warehouse in an economical and efficient way.

3.4.1 Storage tiers and tier redundancy

SC Series arrays support up to three storage tiers. With the advent of flash drives, it is common to have a two-tier system that consists of SSD drives in tier 1 and HDDs in tier 3. Further optimization and protection are done within each tier with different RAID levels.

The storage tiers and tier redundancy are automatically configured by SCOS based on the drive types and sizes. Adjustment to the storage tiers and tier redundancy can be done in DSM at any time. However, it is recommended to make the adjustment only after careful review and evaluation.

Table 1 shows an example of a simple two-tier SC Series storage configuration for a data warehouse. The capacity tier is ideal for storing a large amount of historical data while the high performance tier is great for handling the occasional transactions and queries.

Table 1 A two-tier SC Series storage configuration

Tier type	Storage tier/RAID level	Media
High Performance Tier	T1/R10	SSDs
	T1/R5 or R6	
Large Capacity Tier	T3/R10	15K, 10K, or 7K HDDs
	T3/R5 or R6	

Note: The document [Understanding RAID with Dell SC Series Storage](#) provides insights on how Dell EMC implements different RAIDs and storage tiers in an SC Series array.

3.4.2 Dynamic data placement with storage profiles

It is important to understand the different types of data stored in a typical data warehouse and how these data are used. A typical data warehouse stores three types of data:

- Historical data
 - Requires a large amount of storage and long retention time
 - Constantly grows over time
 - Data is fairly static
 - Is read to generate summarized or calculated data
 - Updates or new data are added through batch jobs
- Summarized or calculated data
 - New data is created by the Extract, Transform, and Load (ETL) process
 - Reports are generated by on-demand or scheduled queries
 - Data gets refreshed periodically
 - Typical reporting and user query response time ranges from several seconds to minutes
- Metadata
 - Provides information about other data and schema objects within a database
 - Is used by applications to locate and compute data correctly
 - Is fairly static until refresh

Data stored in an SC Series array is not permanently tied to a specific storage tier, type, or RAID level. Each volume has a selected **storage profile**. The profile dictates which storage tier accepts initial writes and how Data Progression moves data between tiers and RAID levels to balance performance and storage efficiency.

Note: Storage profiles are predefined based on the available drive types in the array. It is important to review the storage profiles information in appendix A and in the *Dell Storage Manager Administrator's Guide*.

The following list shows an example of how to apply different storage profiles to different data types:

- Dell EMC and Oracle recommend using Oracle ASM to manage the data storage.
- The same storage profile must be applied to each SC Series volume that constitutes the ASM disk groups.
- In most cases, the predefined **Recommended** profile offers the most effective way to manage data across all storage tiers and RAID levels. The Recommended profile dictates that new data writes to tier 1 R10, moves snapshot data to R5/R6 within the same tier, and moves aged or inaccessible data to tier 3.
- Alternatively, certain data, such as summarized data and historical data, can be restricted to storage tier(s) by assigning the tier specific storage profiles such as **High Priority** or **Low Priority**. Data can still move between different RAID levels within the same tier to gain efficiency. This offers administrators more precise control over the data placement when appropriate.
- Storage profiles can be set or changed in the **Edit Volume** dialog box in DSM at any time (see Figure 3). Therefore, administrators can experiment with the different storage profiles to find the best fit for their needs. Data Progression moves data into appropriate tiers during the next Data Progression cycle without impacting running applications.

Table 2 offers an example of selected storage profiles on the three data warehouse data types.

Table 2 Data warehouse data types and storage profiles

ASM disk groups	Data type	Storage profiles	Storage tiers/RAID levels	Description
+DATA1DG	Historical data	Recommended	All tiers/ all RAID levels	<ul style="list-style-type: none"> • Consumes large amount of storage • Predictable and consistent growth • Data is mostly read-only • New data is written to the fastest write layer T1 R10 • Active data remains in T1 • Aged data moves to T3 R5 or R6 which offers the same read performance as R10 but takes less space
		Low Priority (Tier 3)	T3/ all RAID levels	<ul style="list-style-type: none"> • New data writes to the fastest write layer T3 R10 • Active data remains in T3 R10 • Aged data moves to T3 R5 or R6 which offers as good read performance as R10 with less space
+DATA2DG	Summarized/ calculated data	Recommended	All tiers/ all RAID levels	<ul style="list-style-type: none"> • New data writes to the fastest write layer T1 R10 • Active data remains on T1 • Aged data moves to T3 R5 or T6 which offers as good read performance as R10 with less space

		High Priority (Tier 1)	T1/ all RAID levels	<ul style="list-style-type: none"> Pin data in T1 only if it is deemed highly critical and requires the best performance New data writes to the fastest write layer T1 R10 Aged data can still move to T1 R5 or R6 to free up space
+SYSTEMDG	Metadata Redo logs System data	Recommended	All tiers/ all RAID levels	<ul style="list-style-type: none"> Metadata changes infrequently; mostly read I/Os Redo logs are circular logs; mostly write I/Os New data/logs writes to the fastest layer T1 R10 Aged data moves to T3 R5 or R6
+FRADG	Archived logs	Low Priority (Tier 3)	T3/ all RAID levels	<ul style="list-style-type: none"> New logs writes to T3 R10 Aged logs moves to T3 R5 or R6 to free up space

For LVM, the three main data types are:

- Oracle software binaries
- Log files generated by various Oracle processes
- Dump files generated by Oracle Data Pump utility

The binaries are read into the system memory during the startup of the environment. They do not get accessed on storage or changed frequently. For the log files, they are continuously written to the file system but only get read when administrators need to review and troubleshoot an issue. Therefore storing these data types on tier 3 storage would seem reasonably adequate.

Table 3 LVM data types and storage profiles

LVM disk group	Data type	Storage profiles	Storage tiers/RAID levels	Description
vgoracle	<ul style="list-style-type: none"> • Oracle software • Oracle logs • Oracle data pump files 	Low Priority (Tier 3)	T3/All RAID levels	<ul style="list-style-type: none"> • New files writes to T3 R10 • Aged files move to T3 R5 or R6 to free up space • T3 R5/R6 offers good read performance for logs and dump files

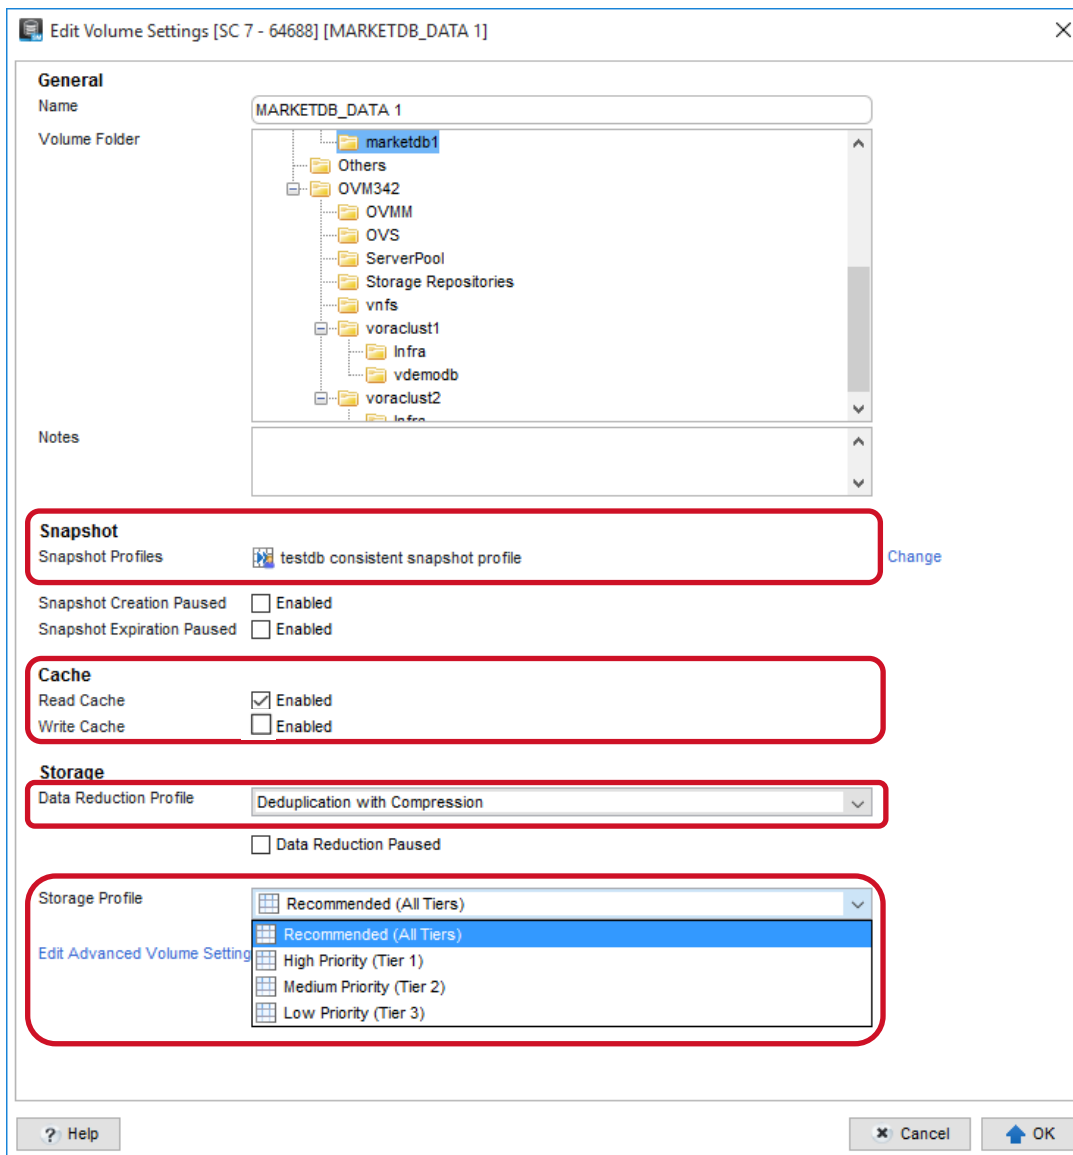


Figure 3 Volume attributes in DSM

3.4.3 Initial loading of historical data

If a large amount of data will be transferred into the data warehouse initially, it might not be desirable to consume all tier 1 storage for the initial load. A common practice is to direct data during ingestion to the lowest tier where capacity is abundant. In DSM, select the **Low Priority** storage profile on each volume for the duration of the ingestion. After the initial loading is complete, change the profile to **Recommended**. This allows Data Progression to start relocating the data to different RAID levels and tiers based on actual I/O patterns in the subsequent progression cycles.

3.4.4 Data reduction

In addition to thin-provisioning and multi-RAID configuration, the SC Series deduplication and compression further reduce the space consumption on the storage array. When deduplication is enabled on a volume, it automatically enables compression as well, though compression can be enabled without deduplication. The space savings is largely dependent on the data stored in a data warehouse. Since the features can be turned on or off at any time, administrators are encouraged to enable the features and assess their performance and

space-saving benefits at the most convenient time. The following offers an example where data reduction features might be suitable in a data warehouse.

Note: Do not enable SC Series data reduction on the same set of data where database/application-level data reduction is already enabled. There might not be significant reduction savings on already reduced data.

Table 4 Data reduction model based on data types

ASM disk groups	Data type	Data reduction
+DATA1DG	<ul style="list-style-type: none"> Historical data 	Deduplication with compression
+DATA2DG	<ul style="list-style-type: none"> Customized/summarized data 	None
+SYSTEMDG	<ul style="list-style-type: none"> Metadata Redo logs System data Temp data 	None
+FRADG	<ul style="list-style-type: none"> Archived logs 	
LVM disk groups	Data type	Data reduction
vgoracle	<ul style="list-style-type: none"> Oracle binaries Oracle logs Oracle data pump files 	Deduplication with compression ¹

¹ Internal testing shows 40–50% savings on Oracle binaries, logs, and data pump files when deduplication and compression are turned on. The actual savings will vary based on the amount of files accumulated in the file system.

Data Reduction is set in the **Edit Volume** dialog box in DSM (see Figure 3). Three options are available:

- Deduplication with Compression
- Compression
- None (Default)

It is recommended to enable the feature on the database volumes for a period of time long enough to go through a typical workload cycle. The amount of space savings is reflected in DSM after Data Progression has run. The information can be found in the following locations within DSM:

- For volume-specific data-reduction statistics, the savings are displayed under the volume **Statistics** tab. See Figure 4.
- For array-wide data-reduction summary statistics, the savings are displayed under **Storage Types > Summary** tab. See Figure 5.
- To see the data-reduction savings of all volumes in one place, go to **Storage Types > Volumes** tab. See Figure 6.

For more information on SC Series data reduction, refer to the *Dell Storage Center OS 7.0 Data Reduction with Deduplication and Compression Solution Guide*.

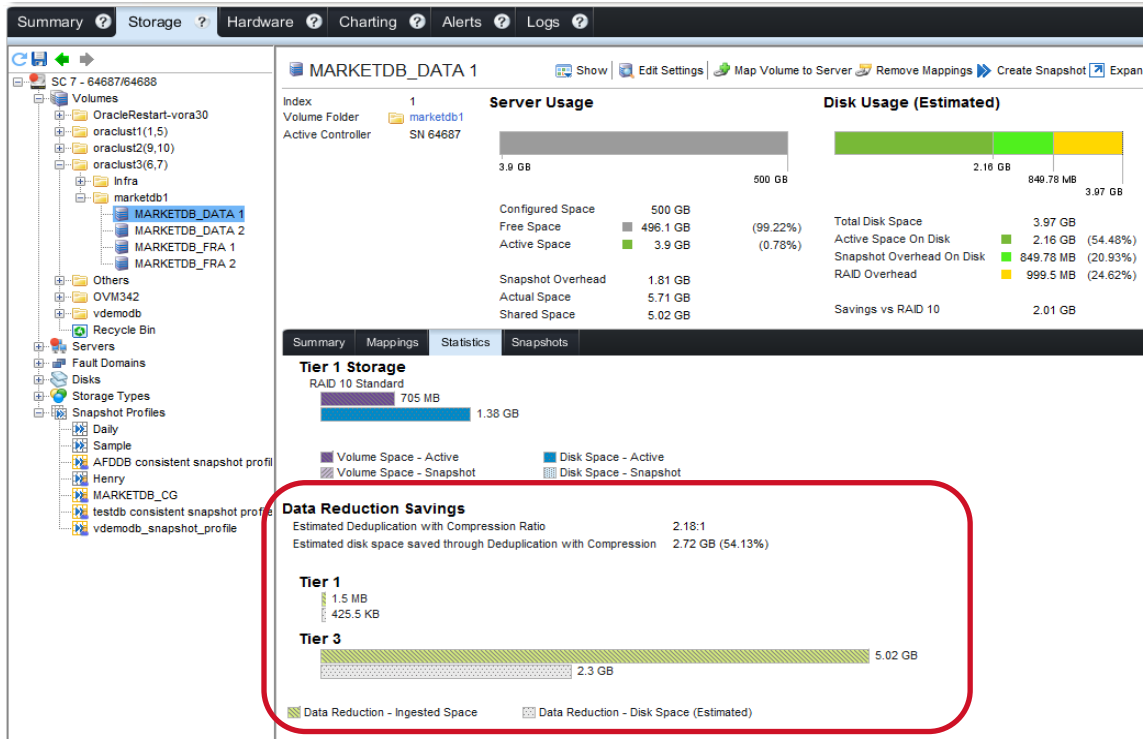


Figure 4 Per-volume data-reduction statistics

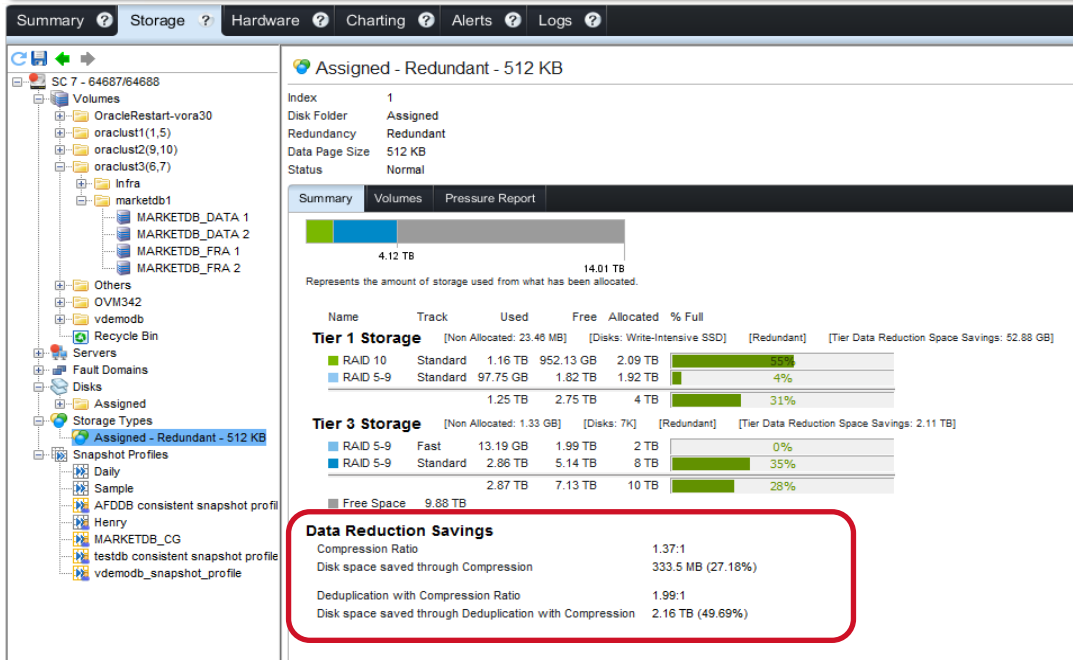


Figure 5 Array-wide data-reduction statistics

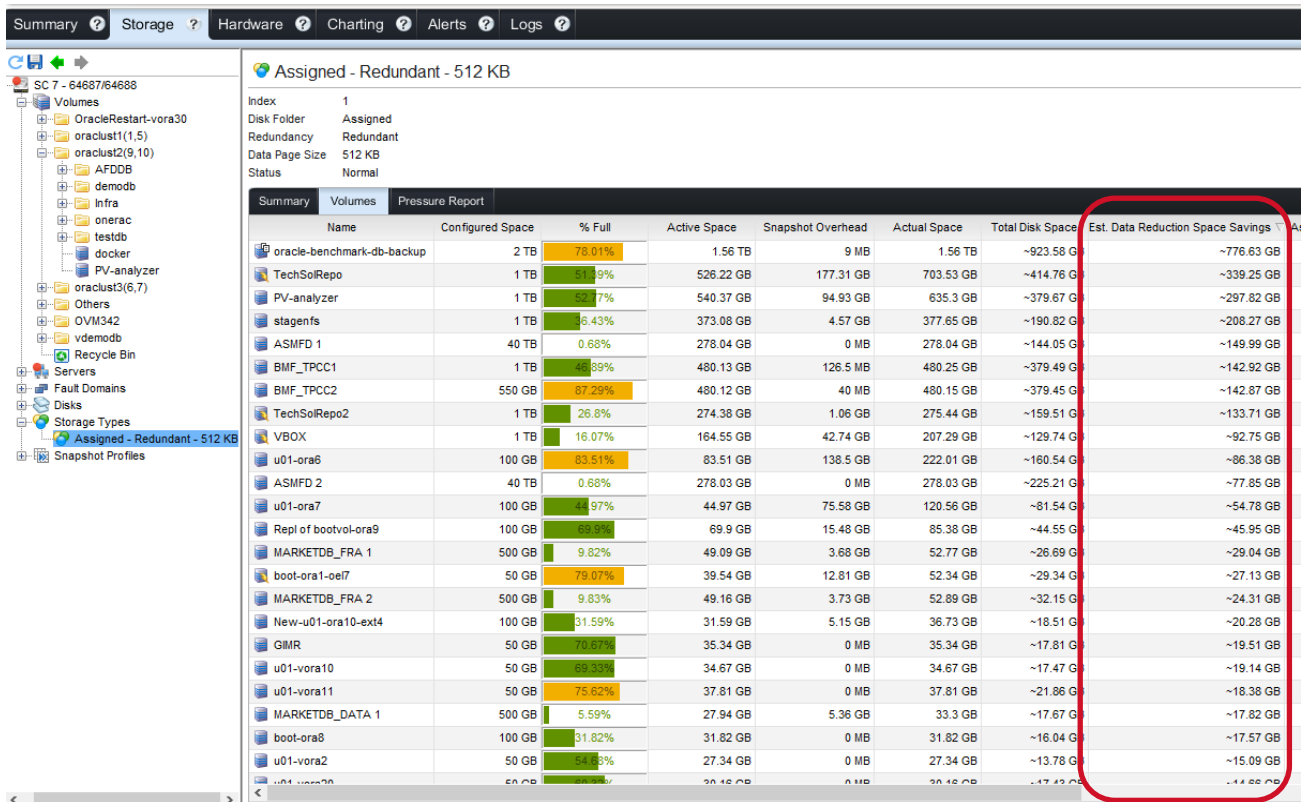


Figure 6 Volume space allocation statistics

3.4.5 Thin provisioning compared to preallocating storage

By default, all SC Series volumes are thin provisioned (the SC Series array allocates pages to the volume as data is written to it). Alternatively, the whole volume can be preallocated at the time of creation. However, regardless of the actual space consumption, the full volume size is reserved as used and cannot be reassigned to other volumes.

Note: In most use cases, Dell EMC recommends thin provisioning volumes because it allows the array to adapt to the actual space consumption.

Administrators can override this behavior on a per-volume basis. To allow preallocation of volume space, in DSM, edit the SC Series settings, and under **Preferences**, check **Allow Preallocate Storage Selection**. When creating a volume, a server must be selected before the **Preallocate Storage** option becomes available.

Note: The preallocation operation is irreversible. To release the space, create a new thin provisioned volume, migrate the data over, and delete the previously preallocated volume.

To realize the benefit of these thin-provisioned volumes, avoid preallocating storage on Oracle data files.

When creating the tablespaces and data files on these ASM disk groups, administrators can set an initial size of each data file and specify the autoextend clause to include an autoextend size on each data file. This allows SC Series array to allocate storage for the initial data file size. When the data files are full, additional space will be allocated by the amount of autoextend size.

For example, the following SQL statement creates a bigfile tablespace intended for storing the historical data. The statement preallocates 1TB of space initially. When it is full, additional storage will be added in 1GB amounts.

```
create bigfile tablespace HISTDATATS datafile '+DATA1DG' size 1024G autoextend on next 1024M maxsize unlimited;
```

If a large amount of data will be loaded into the tablespace initially, preallocate enough storage to handle the load and to improve the loading speed.

Choose an appropriate autoextend size to minimize the frequency of extending the data files and allocating storage in the SC Series array during the regular data ingestion cycle. This provides the benefit of thin provisioning and gradually allocating storage as needed without incurring excessive overhead.

3.4.6 Volume cache

Read and write cache generally improve performance for an all-HDD configuration. However, when SSDs are used in the SC Series array, it is recommended to disable the write cache on volumes that use SSD storage to improve performance. This means that any volume with a storage profile that writes to T1 storage such as **Recommended** and **High Priority** should disable the write cache.

To enable or disable the read or write cache, in DSM > SC user preferences, enable the **Allow Cache Selection**. See the *Dell Storage Manager Administrator's Guide* for instructions.

3.5 Protecting Oracle with snapshots and a consistency group

Due to the large data volume in a data warehouses, backing up or restoring such an environment might take days to complete, hence, affecting the organization's ability to meet a specific recovery point objective (RPO) and recovery time objective (RTO).

The SC Series storage snapshot feature is an ideal solution to address these challenges. Enabled through the **Snapshot Profile** on each volume (see Figure 3), it protects an Oracle database by creating fast and space-efficient snapshots of the storage volumes. A snapshot profile is a policy-based collection of rules describing the type of schedule (once, daily, weekly, or monthly), a date/time and the interval to create the snapshot, the volumes, an expiration time for the snapshot, and if the snapshot should be write-consistent across all volumes in the schedule.

Guidelines for SC Series snapshots:

- Create a dedicated snapshot profile in DSM for each data warehouse/database.
- When a snapshot is taken, existing pages are frozen and no data is moved or copied. New data or changed data are then written to new pages. Taking a complete snapshot of a data warehouse usually takes seconds, regardless of the amount of data involved.
- Each complete set of snapshots represents a full point-in-time copy of the data warehouse.
- The snapshot size will depend on the amount of changes to the data since the last snapshot was taken. Generally, it is significantly less than the original full size database.
- Thousands of snapshots are supported per volume.
- Data Progression moves frozen pages according to the selected storage profile to optimize storage. See appendix A for storage profiles information.
- Snapshots offload the traditional backup processing, such as RMAN, from the primary server.
- They simplify the cloning of a full-size data warehouse for testing, support, or departmental use.

- Snapshots allow fast recovery and improve RTO and RPO. Data from snapshots are immediately available through view volumes.
- An instant backup should be created before major upgrades or changes in the data warehouse.

Considerations for planning snapshot schedules:

- Snapshots can be taken on a schedule or on-demand.
- Schedule snapshots based on backup policy and RPO. SC Series storage allows a schedule as frequent as every 5 minutes.
- Schedule based on the data ingestion cycle. Take a snapshot after each mass data load. This allows Data Progression to optimize the frozen pages by moving to different RAID levels and storage tiers.
- Set an expiration policy to automatically reclaim storage space from obsolete snapshots.
- Multiple schedules can be attached to the same snapshot profile. It is highly recommended to have at least one daily snapshot scheduled for every SC Series volume.

When a database spans across multiple storage volumes like the example shown in Figure 2, the storage volumes of the same database must be assigned to the same snapshot profile and the **consistent** method must be chosen for the profile. Otherwise, Oracle might not be able to recover from a multi-volume snapshot. This **consistency group** feature allows taking a database-consistent snapshot that spans across multiple volumes. On the Oracle side, use **ALTER DATABASE BEGIN BACKUP** and **END BACKUP** before and after the snapshot is taken. Similarly, a consistency group should be configured for multi-volume LVM.

To access or restore data from a snapshot, a **view volume** is created and mounted with full read-write access on any recovery system in DSM. A view volume shares the frozen pages with the parent volume but data written to the view volume is tracked separately and consumes real storage space. A view volume is treated like a regular volume in that it has its own volume attributes and changes to the view volume does not affect its parent volume.

3.6 Sizing considerations

When sizing an SC Series array for a data warehouse, the goal is to maximize I/O bandwidth by having enough physical disks in each tier and channels accessing the data. It is important to work closely with your Dell EMC account team to collect and understand the information described in the following sections.

3.6.1 Collect and measure performance statistics

Performance expectations must be defined by quantifiable measurements and clearly communicated to all teams involved, including the Dell EMC team, so that they can properly size the SC Series array to meet the objectives. Collect performance statistics in an existing environment that has similar workload and activities to the final environment is strongly preferred.

Some operating systems provide native tools that collect basic IOPS information. For example, on Linux, **sar** and **iostat** come pre-installed on the system. Other free utilities like **atop** and **collectl** are similar to sar but have additional features. On Microsoft Windows, **Performance Monitor** is readily available. Dell EMC also offers a free tool, the **Dell Performance Analysis Collection Kit (DPACK)**, to assist in collecting and analyzing system and I/O performance data. For information on DPACK, see the [DPACK FAQ](#) or contact the Dell EMC team.

For an Oracle database, the utilities used most are **statspack** and **AWR**. AWR is preferred to the older statspack, but either one can provide abundant performance statistics of a database.

3.6.2 Configure I/O for bandwidth

Storage configurations for a data warehouse should be chosen based on the I/O bandwidth and not just on the overall storage capacity. The following factors should be considered:

- The expected throughput of the data warehouse and all other applications that share the SC Series array.
- The number of drives — more drives increase overall bandwidth
- Types of drives — SSDs offer higher transfer bandwidth per drive. Therefore, fewer SSDs in Tier 1 might be used to provide similar bandwidth as HDDs.
- Higher speed 12Gb SAS channels are available for back-end drive enclosure connectivity in some SC Series arrays. Choose the highest speed wherever available.
- Account for all components along the I/O path and ensure that they are adequately configured to handle the total aggregated bandwidth, for example, the array front-end FC/iSCSI/SAS channels, the number and speed of server HBAs, and connectivity.

3.6.3 Capacity planning

The advanced features of SC Series storage such as thin provisioning, auto-tiering, multi-RAID levels, data reduction features, and the snapshot capability make the best and most efficient use of storage space possible. To figure out the capacity needed to support a data warehouse, consider the following factors that have direct effects on the storage consumption. Also, storage capacity can be dynamically added or removed from the SC Series array without affecting the applications. Thus, customers only invest in what they need and still maintain the flexibility to meet any future growth. When new drives are added, data is redistributed and rebalanced across all drives automatically to maximize performance.

Consider the following when planning for capacity needs:

- RAID types, such as RAID 10 or RAID 5/6
- Use of deduplication and compression
- Snapshot frequency and retention policy
- Number of data warehouse environments, full-size copy plus thin clones
- Data growth rate, factoring new data added each day, plus new archive logs, plus new customized/summarized data
- Data change rate, considering updates or refreshes of existing data and reports
- Oracle ASM redundancy level: 1x size for External, 2x size for Normal, and 3x size for High
- Storage profiles
- Amount of archive logs to keep
- Tier 1 storage should have enough storage to hold the active data and aging data until they are aged out to lower tier. For accessible data pages, the aging period is 12 Data Progression cycles. For inaccessible data pages, they are migrated to the lower tier in the next Data Progression cycle. See the *Dell Storage Manager Administrator's Guide* and *Dell EMC SC Series Arrays and Oracle* for more information on snapshots and Data Progression.
- Estimate storage per tier. For example:
 - Estimate tier 3 storage based on the amount of historical data, plus snapshots, plus data growth.
 - Estimate tier 1 storage based on the amount of customized/summarized data, plus frequently accessed data, plus snapshots, plus new data.

Oracle Automatic Workload Repository (AWR) can be used to help estimate how much storage is required for snapshots. It provides detailed usage information on redo logs, archived logs, temporary space, and data files. Administrators should be able to deduce the data-change rate from the AWR reports. Reports should cover different time periods to coincide with the time the snapshots are intended to be taken, and the information used as the basis of the calculation.

Dell EMC is well equipped to help customers analyze both storage and performance requirements based on many of these inputs.

3.6.4 Test the I/O system before implementation

I/O performance should be validated and tested before putting the environment into full production mode. Tools such as Oracle ORION, vdbench, and FIO provide a comprehensive set of features to exercise and measure the storage. ORION has an advantage over others because it is explicitly designed to simulate Oracle database I/O workloads using the same I/O software stack as Oracle. It also provides a data warehouse simulation mode which simplifies the setup and execution of the test. ORION has been bundled with the Oracle database software and can be found in \$ORACLE_HOME/bin directory. See appendix B.1 for references to these performance tools. For more information on how to configure and run ORION, see chapter, [Calibration with the Oracle Orion Calibration Tool](#), in Oracle Performance Guide.

4 Conclusion

SC Series arrays provide well-balanced storage solutions that benefit large data warehouses by reducing the total cost without sacrificing the performance. Thin provisioning, continuous self-optimization, and data reduction work seamlessly together to provide the most cost-efficient solution. Performance is delivered where it is needed through auto-tiering and multi-RAID levels. When the data warehouse grows in the future, the array can scale easily by adding drives at any time. Data is automatically rebalanced across all available drives.

In summary, when deploying an Oracle data warehouse on SC Series storage, use the recommendations outlined in this document for improved manageability and performance:

- Configure I/O for bandwidth and not just capacity.
- SC Series volumes for ASM should follow the Oracle recommendation on LUN sizing.
- Disable write cache on SC Series volumes that write to the SSD tier.
- Select storage profiles that benefit different data warehouse data types.
- Enable data reduction to further reduce space consumption.
- Use snapshots to create space-efficient copies of a data warehouse and improve RPO and RTO.
- Use multipath software.
- Use Oracle ASM and Oracle managed files.
- Use ASM external redundancy, unless business needs require normal or high redundancy.
- Make sure ASM disk groups contain multiple disks to sustain striping far and wide and to reduce Linux kernel contention accessing and queuing for the same disk.
- Perform due diligence on I/O system tests before database implementation.
- Assign all volumes used by a specific database to the same consistent snapshot profile.

A Storage profiles and RAID levels

Table 5 Storage profiles available with write SSDs and HDDs, or with HDDs only

Name	Initial write tier	Tier (T) and RAID levels	Progression
Recommended (All Tiers)	1	Writes: T1 RAID 10 Snapshots: RAID 5/RAID 6	To all tiers
High Priority (Tier 1)	1	Writes: T1 RAID 10 Snapshots: T1 RAID 5/RAID 6	Stay in tier 1
Medium Priority (Tier 2)	2	Writes: T2 RAID 10 Snapshots: T2 RAID 5/RAID 6	Stay in tier 2
Low Priority (Tier 3)	3	Writes: T3 RAID 10 Snapshots: T3 RAID 5/RAID 6	Stay in tier 3

Table 6 Storage profiles available with both write SSDs and read SSDs

Name	Initial write tier	Tier (T) and RAID levels	Progression
Flash Optimized with Progression (Tier 1 to All Tiers)	1	Writes: T1 RAID 10 Snapshots: RAID 5/RAID 6	To all tiers
Write Intensive (Tier 1)	1	Writes: T1 RAID 10 Snapshots: T1 RAID 10	Stay in tier 1
Flash Only with Progression (Tier 1 to Tier 2)	1	Writes: T1 RAID 10 Snapshots: T2 RAID 5	To tier 2
Low Priority with Progression (Tier 3 to Tier 2)	3	Writes: T3 RAID 10 Snapshots: T3 RAID 5/6 or T2 RAID 5	To tier 2
Low Priority (Tier 3)	3	Writes: T3 RAID 10 Snapshots: T2/T3 RAID 5/6	Stay in tier 3

Table 7 RAID stripe width in SC Series arrays

RAID Level	Stripe width description
RAID 0	Stripes across all drives in the tier with no redundancy
RAID 10	Stripes data along with one copy across all drives in the tier
RAID 10-DM	Stripes data along with two copies across all drives in the tier
RAID 5-5	Distributes parity across five drives (4 data segments, 1 parity segment for each stripe); tier requires at least 6 drives (5 for RAID, and one for spare)
RAID 5-9	Distributes parity across nine drives (8 data segments, 1 parity segment for each stripe); tier requires at least 10 drives (9 for RAID, and one for spare)
RAID 6-6	Distributes parity across six drives (4 data segments, 2 parity segments for each stripe); tier requires at least 7 drives (6 for RAID, and one for spare)
RAID 6-10	Distributes parity across ten drives (8 data segments, 2 parity segments for each stripe); tier requires at least 11 drives (10 for RAID, and one for spare)

B Technical support and resources

[Dell.com/support](https://dell.com/support) is focused on meeting customer needs with proven services and support.

[Storage technical documents and videos](#) provide expertise that helps to ensure customer success on Dell EMC storage platforms.

B.1 Additional resources

Referenced or recommended Dell EMC publications:

- *Dell Storage Manager Administrator's Guide* available on the [Dell Support website](#).
- Dell Storage Center OS 7.0 Data Reduction with Deduplication and Compression
- Understanding RAID with Dell SC Series Storage
- Dell SC Series Storage with Oracle ASM Best Practices
- Dell Storage Center with Red Hat Enterprise Linux (RHEL) 6x Best Practices
- Dell Storage Center with Red Hat Enterprise Linux (RHEL) 7x Best Practices
- Oracle Backup and Recovery Best Practices for Dell SC Series Storage
- Oracle OLTP with Dell Storage SC4020 All-flash Optimized Arrays
- Dell Storage Center Replication and Oracle Data Guard
- Dell EMC SC Series Arrays and Oracle
- Dell SupportAssist for SC Series Arrays
- Dell Storage PowerShell SDK Cookbook
- Dell SC Series Storage: Synchronous Replication and Live Volume

Referenced or recommended Red Hat publications:

- [Red Hat Enterprise Linux 6 DM Multipath, DM Multipath Configuration and Administration](#)

Referenced or recommended Oracle publications:

- [Oracle Automatic Storage Management Administrator's Guide](#)
- [Oracle Database Administration Documentation Library](#)
- [Oracle Performance Guide](#)

Referenced or recommended I/O benchmark resources:

- [vdbench download and documentation](#)
- [FIO download](#)
- [FIO output explained](#)