

Comprehending the Tradeoffs between Deploying Oracle® Database on RAID 5 and RAID 10 Storage Configurations

A Dell® Technical White Paper

Database Solutions Engineering

By Sudhansu Sekhar and Raghunatha M

Dell Product Group

April 2009



Executive Summary

Redundant Array of Independent Disks (RAID) technology is used to improve performance and availability of the underlying storage arrays for any application. With RAID implementations, performance is improved by distributing I/O across many physical disks and availability is increased with the data rebuilding capability. But both high availability and high performance comes at a price, which are different for different RAID implementations.

Most of the time database and storage administrators ask the age-old question—whether to go with RAID 10 or RAID 5 for their database deployment. Customers struggle to make right tradeoffs between these two RAID implementations for their database deployments. This white paper evaluates different tradeoffs including performance, cost, and availability while evaluating RAID 10 compared to RAID 5. This will help the customers to make informed decisions to choose the right RAID implementation for their specific workload which can meet their desired performance and availability requirements.

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© 2009 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

Dell, the *DELL* logo, and the *DELL* badge are trademarks of Dell Inc. *EMC* is a registered trademark of EMC in the United States and/or other countries. *Oracle* and *RAC* are registered trademarks of Oracle in the United States and/or other countries.

Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims proprietary interest in the marks and names of others.

Contents

- Executive Summary..... ii
- Introduction 4
- Overview of RAID 4
 - RAID Implementation in Dell®/EMC® Storage Arrays 5
- RAID Tradeoffs 5
- Performance Tradeoff — RAID 10 Compared to RAID 5 6
 - Test Configuration..... 6
 - Test Cases..... 6
 - OLTP Workload 6
 - Results and Analysis 7
 - OLAP Workload 9
 - Results and Analysis 9
- Cost Tradeoff — RAID 10 Compared to RAID 5 11
 - Test Configurations 11
 - Test Cases..... 12
 - Results and Analysis 12
- Availability Tradeoff – RAID 10 compared to RAID 5 14
 - Test Configurations 14
 - Test Cases..... 14
 - Results and Analysis 14
- Four Quadrant Analysis..... 15
- Conclusion..... 18
- GLOSSARY..... 19
- References 21

Introduction

Configuration of the underlying storage subsystem is one of the most critical tasks while deploying applications. For I/O intensive applications like Oracle Databases, the storage configuration can be extremely complex. Storage and database administrators need to comprehend several performance criteria while considering the storage design. Some of the important storage options that need careful attention are RAID configurations, cache configurations, disk types, and so on. Of these, one topic of debate is the RAID configuration for storage arrays.

From a performance perspective, RAID 10 is claimed to be the ideal RAID choice for database deployments. However, customers sometimes have to make cost tradeoffs, which force them away from RAID 10 and lead them to choose the more cost effective RAID 5 configuration.

Both RAID 10 and RAID 5 configurations have advantages and disadvantages and are suitable for specific I/O patterns. This white paper provides recommendations on using RAID 10 and RAID 5 for different database workloads. This paper also examines several tradeoffs (for example, performance tradeoff, cost tradeoff, and availability tradeoff) while deploying Oracle Databases with RAID 10 compared to RAID 5. A four-quadrant analysis was provided to help customers to choose the RAID configurations for databases with mixed-use deployment, where customers run online transaction processing (OLTP) during the day and DSS workload at night. This will help the database and storage administrators to make an informed decision when planning their complex database deployments.

Oracle Database has several components and each of these components can have different I/O patterns. Components like online redo logs, control files, temp and undo table spaces are write-intensive. On the other hand, data files for Oracle Database can be read- or write-intensive, depending on the type of application which is accessing the database. As a result, different RAID technologies can be appropriate for different database components. Customers can even mix and match configurations for their database deployment by placing different components of Oracle Database in different RAID groups. These types of configurations are analyzed, and recommendations are provided in this white paper.

Overview of RAID

The objective of the RAID technology is to accomplish the following goals:

1. Improve the disk read/write performance (by parallelizing the I/O over multiple disks)
2. Improve redundancy (by introducing data recovery capability in the event of disk failure)

Several RAID technologies are available to achieve better disk performance and data redundancy. Below are the most commonly used RAID configurations used for the implementation of Oracle Database:

- RAID 10 (also known as stripe and mirror) — This technique is a mix of RAID 1 (which is mirroring) and RAID 0 (which is striping). With RAID 10, data is spread across multiple disks (to

Using RAID 10 Compared to RAID 5 Configurations for Oracle® Database Deployments — Tradeoffs

improve performance), and the striped data is mirrored as a different dataset (to provide redundancy).

- RAID 5 — The implementation of a mathematical exclusive or (XOR) is a method that uses parity to reconstruct data from a failed drive. RAID 5 does not mirror the whole set of data to build redundancy into the setup. Data is striped across all the disks to also improve the performance in RAID 5.

RAID Implementation in Dell®/EMC® Storage Arrays

Different storage vendors implemented the RAID technologies in several ways. Below is the description on the way RAID 10 and RAID 5 are implemented in Dell/EMC storage arrays.

For Dell/EMC storage arrays, RAID 10 takes care of both data striping and mirroring. First, the data is striped across all the disks in the RAID group. This ensures the optimum performance for the RAID configuration. Once the data is striped across all the disks, the data sets get mirrored inside the available disks to create the redundant copies.

For optimal redundancy, EMC recommends to configure the RAID 5 with a 3+1 configuration. With a 3+1 configuration, one disk is used for storing parity data for every three disks that are used in the RAID configuration. For storage systems, RAID 5 can be implemented with 3- to 16-disk RAID groups. But considerations should be taken while selecting the number of disks for the RAID group. The larger the size of the RAID group, the more time it takes to rebuild the data in the event of a disk failure.

RAID Tradeoffs

As discussed in the previous section, RAID implementations improve performance and redundancy. However, performance improvement and redundancy come with cost as a larger amount of disks is required, for example:

- For the same size of database, RAID 10 requires more disks than RAID 5 to achieve the desired data availability.
- RAID 5 requires more time to reconstruct the failed data in case of disk failure.
- RAID 5 has to deal with parity overhead for disk writes.

Customers need to consider the following tradeoffs before implementing a particular RAID technology:

- **Performance tradeoff** — the performance benefit of preferring RAID 10 to RAID 5 with the desired I/O pattern
- **Cost tradeoff** — the cost benefit of either RAID 10 or RAID 5 implementation
- **Availability tradeoff** — the amount of time required to recover data from a failed disk scenario

The following sections discuss the above considerations in the context of Oracle Database deployed on RAID 10 compare to RAID 5 configurations.

Performance Tradeoff — RAID 10 Compared to RAID 5

This section provides the performance analysis for both RAID 10 and RAID 5 configurations for different database workloads and I/O patterns. The analysis can be used for determining the performance tradeoffs between RAID 10 and RAID 5 implementations for different database workloads.

Test Configuration

For the performance tradeoff test, two Oracle 11G R1 (11.1.0.7) two-node RAC clusters were deployed. One RAC cluster was deployed with RAID 10 layout and the other cluster with RAID 5 layout. Both of the RAID 10 and RAID 5 was created from eight disks. The following hardware and software components were used for these two clusters:

Hardware	Components	RAID 10 and RAID 5
	Server	2xR710
	Storage	CX4-480
	FC switch	Brocade SW 5000
	Network switch	Dell™PowerConnect™ 5448
Software		
	Operating system	RHEL 5.2
	Oracle RDBMS (with RAC)	EE 11G R1 (11.1.0.7)
	EMC PowerPath	V 5.1.0
Database		
	Database size	400 GB
Benchmark Tools		
	Quest Benchmark Factory (OLTP workloads)	V 5.5.1
	ORION (OLAP workload)	10.2

Test Cases

The test cases were divided into two major categories, OLTP workloads and On-Line Analytical Processing (OLAP) workloads:

OLTP Workload

Quest Benchmark Factory has a framework which provides industry standard TPC-C performance runs. Quest Benchmark Factory TPC-C work load is designed for online transaction processing environment with more reads (70%) and less writes (30%). As part of TPC-C performance run, the database was

Using RAID 10 Compared to RAID 5 Configurations for Oracle® Database Deployments — Tradeoffs

loaded with 400 GB of data. Transactions were carried out with increasing user loads. The performance statistics were collected for each of the user loads.

Results and Analysis

For an OLTP workload, performance is typically measured in terms of transactions per second (TPS) and the response time. From a performance prospective, the database should be able to meet the TPS requirement within permissible response time.

Quest Benchmark Factory TPC-C generated the transactions for the database. The transactions per second and the response time metrics are captured for both the RAID 10 and the RAID 5 configurations. Figure 1 provides the TPS with varying user loads, and Figure 2 provides the response time analysis for varying user loads for both the RAID 10 and the RAID 5 configurations.

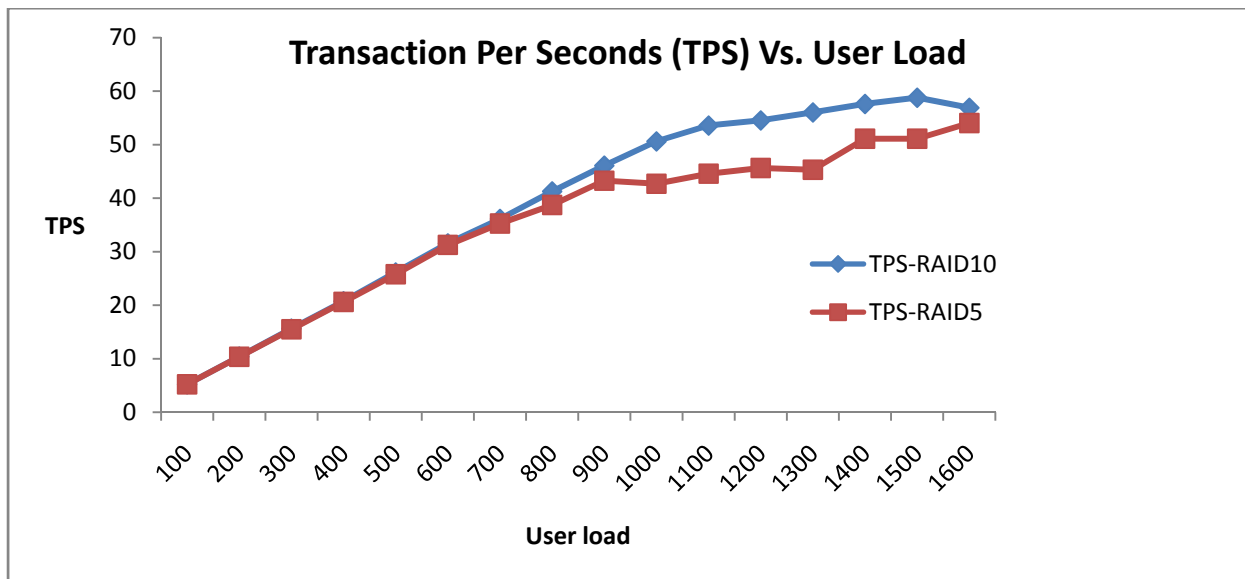


Figure 1: Transactions per Second (TPS) Analysis — RAID 10 Compared to RAID 5

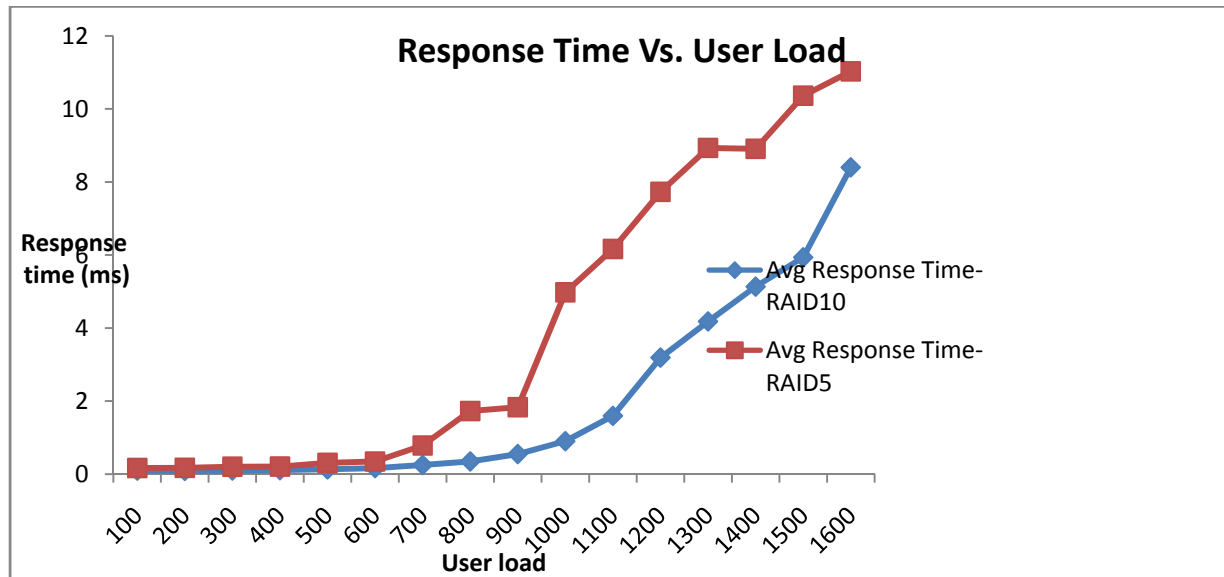


Figure 2: Response Time Analysis — RAID 10 Compared to RAID 5

From Figure 1, we can observe that both the RAID configurations could provide the same TPS around the 900 user load marks. But beyond the load generated by 900 users in this specific benchmark, the RAID 5 cluster is not able to keep pace with RAID 10 cluster.

The response time analysis provides a similar trend for the RAID 10 and the RAID 5 clusters. Response time is the time elapsed for the SQL round trip, in other words, the time elapsed from the time the Benchmark Factory client submits the SQL, to the time at which the client gets the response. From the response time plot (Figure 2), we can observe that the response time increases significantly after the 900 user load mark for RAID 5 cluster. As we can see, the response time exceeds 2 ms beyond 900 user loads for the RAID 5 configuration, while the RAID 10 configuration provides response time below 2 ms up to 1100 user loads.

As a result, for OLTP workloads (loads with roughly a 70/30 read-to-write ratio), the RAID 5 performance can match the RAID 10 configuration up to a certain user load. Beyond that RAID 10 performance is better than that of RAID 5 configuration. In other words, RAID 10 will provide better TPS with less response time compared to RAID 5 configuration beyond certain user load. This behavior is expected because of the parity write penalty associated (for write activities) with RAID 5.

For lightly loaded databases with very low read-and-write workloads, the performance degradation may not be noticed for the RAID 5 configuration. In other words, both RAID 10 and RAID 5 may be used to deploy lightly loaded databases. In this scenario, the parity write penalty will be less because of the light workload.

OLAP Workload

ORION was used to simulate the online analytical processing (OLAP) workload. OLAP workloads often comprise full table scans, long running reports, and backup jobs. Additionally, the I/O size of OLAP is much larger compared to OLTP, mostly in the range of 1 MB. Consequently, the characteristic of the overall data throughput (measured in Megabytes per second) is of interest for these types of workloads.

Databases were stressed with both large sequential read and large sequential write operations. Also, the read-and-write ratio was varied to analyze the throughput behavior with varying queue length.

Results and Analysis

With ORION, the I/O throughputs were captured for ‘large sequential writes’ and ‘large sequential reads’ with 1 MB of I/O size. The throughput was compared for both the RAID 10 and RAID 5 configurations.

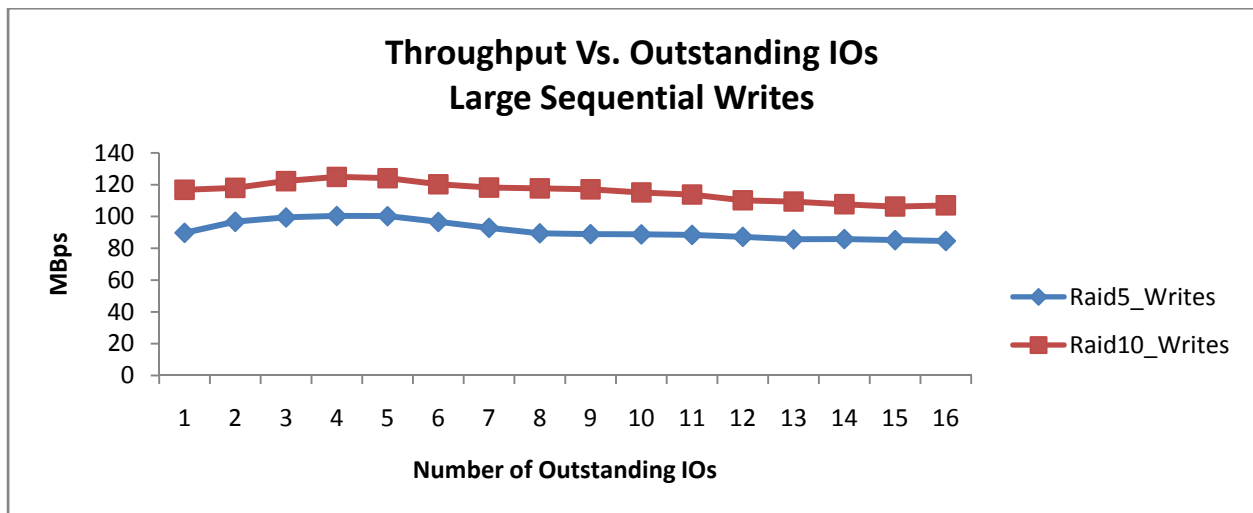


Figure 3: Throughput Analysis for Large Sequential Writes — RAID 10 Compared to Raid 5

Figure 3 provides a comparative analysis of the throughput for large sequential writes for RAID 10 and RAID 5 configurations. The X-axis represents the number of outstanding I/Os and the Y-axis represents the throughput in MBps. As we can observe, the RAID 10 throughput outperforms the RAID 5 configuration for large sequential writes. This behavior is expected because of the parity write penalty associated with the RAID 5 configuration.




Figure 4: Throughput Analysis for Large Sequential Reads — RAID 10 Compared to Raid 5

Figure 4 provides a comparative analysis of the throughput for large sequential reads for RAID 10 and RAID 5 configurations. The X-axis represents the number of outstanding I/Os and Y-axis represents the throughput in MBps. As we can observe, the RAID 5 configuration can match the RAID 10 configuration for large sequential reads. The observed pattern is expected, given that there is no parity write penalty associated with the RAID 5 configuration for read operations.

As mentioned earlier, the read-and-write ratio was varied to analyze the throughput variations with different workloads. For any particular read-and-write ratio (e.g. 10% Read + 90% Write), the throughput was captured both for RAID 10 and RAID 5 configurations. It was observed that for all read-and-write mix, RAID 10 configuration gives a better throughput compared to RAID 5 configuration. The percentage of throughput increase was calculated for each of these read-and-write mixes and plotted as shown in Figure 5. The X-axis has the data point for the different read-and-write mixes (for example, 0R+100W indicates 0% read and 100% write, 10R+90W indicates 10% read and 90% Write, and so on) and the Y-axis has the data point for the percentage of throughput increase for the RAID 10 configuration compared to the RAID 5 configuration.




Figure 5: Throughput analysis for Read-and-write Workload Mix

As we can see from Figure 5, the percentage of throughput gain of RAID 10 over RAID 5 increases with a higher percentage of write operations. The throughput gain for RAID 10 configurations varies from a minimum of 12% (for 0% write) to a maximum of 45% (for 100% write). It is necessary to note that these maximum and minimum values can vary depending upon the read/write cache configuration for the storage arrays. However, the overall recommendation (from the analysis of Figure 5) is that customers may choose to deploy RAID 5 (with minimal performance impact), if their application is mostly read-intensive. In contrast, for every write-intensive application, RAID 10 is recommended for better performance.

Cost Tradeoff — RAID 10 Compared to RAID 5

This section provides the cost benefit analysis for both RAID 10 and RAID 5 implementations. Because RAID 5 implementations provide more usable space compared to RAID 10 implementations, this section assists in the analysis of storage space and to understand the cost benefits while implementing RAID 5.

Test Configurations

Following are the details of the storage array and the disks used for the cost tradeoff analysis.

Hardware	Components	Details
	Storage	CX4-480
	Hard disk	300 GB (RAW space) 265 GB (Usable space)

Using RAID 10 Compared to RAID 5 Configurations for Oracle® Database Deployments — Tradeoffs

Test Cases

To determine the cost tradeoff, two separate RAID groups were created with RAID 10 and RAID 5 layout. For both RAID 10 and RAID 5 RAID groups, the usable space was determined. The percentage of space gained for RAID 10 compared to RAID 5 was calculated for the analysis.

The above steps were repeated for sets of four, eight, and 16 disks, and the trend analysis was done for the cost tradeoff.

Results and Analysis

The following table outlines the details of usable space for RAID 10 and RAID 5 implementations with different sets of disks.

No. of Disks	Disk Size (GB)	Total Capacity (GB)	RAID 5 Usable Space (GB)	RAID 10 Usable Space (GB)
4	300	1073	805	536
8	300	2147	1878	1073
16	300	4294	4025	2146

Table Legend

Disks — Number of disks used in the RAID configuration

Total Capacity — Raw disk capacity (disk capacity × number of disks)

RAID 5 Usable Space — Usable disk space after creating RAID 5

RAID 10 Usable Space — Usable disk space after creating RAID 10

Figure 6 provides the data for the usable space (for both RAID 10 and RAID 5) that the customer can achieve by increasing the number of disks. Figure 7 depicts the trend of increase in usable space for RAID 5 over RAID 10 when increasing the number of disk drives.

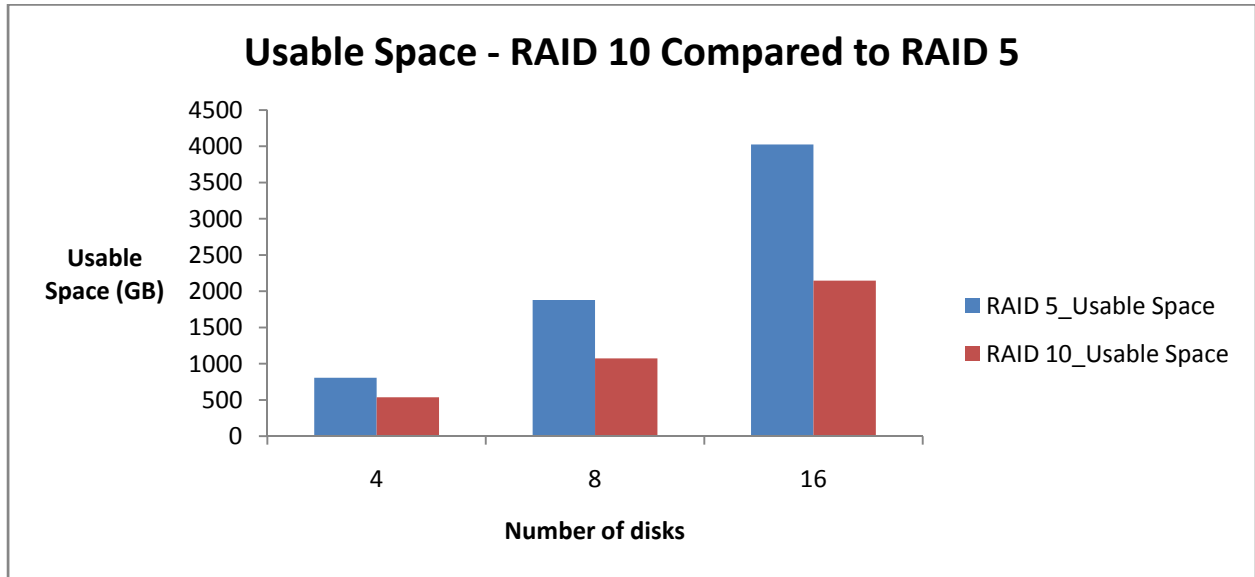


Figure 6: Usable Space Analysis – RAID 10 Compared to RAID 5

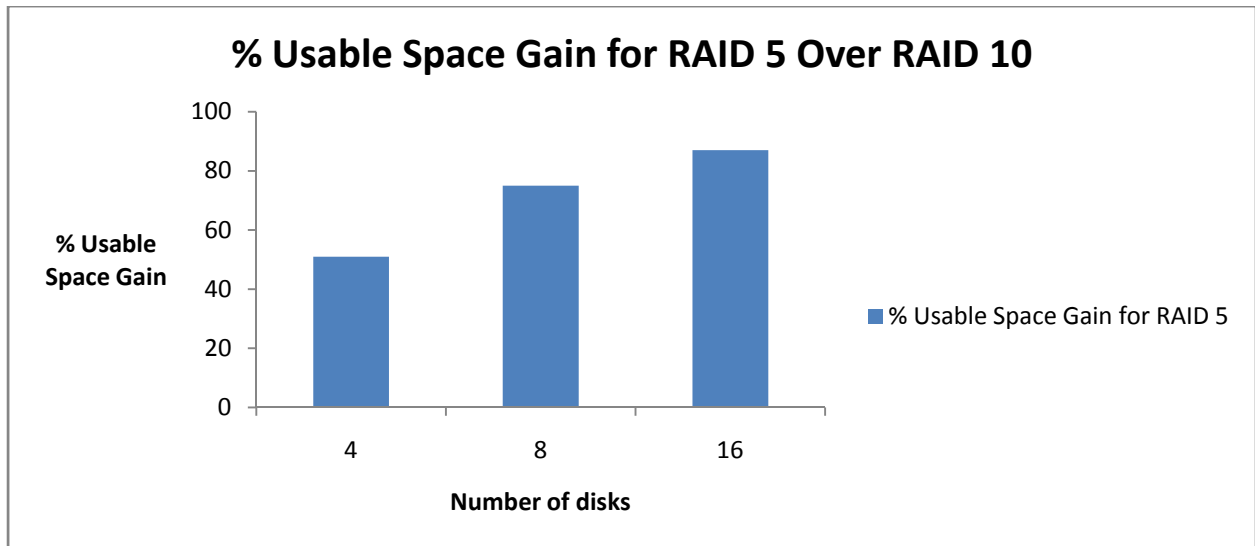


Figure 7: Usable Space Gain for RAID 5 over RAID 10

Comparing Figures 6 and 7, the usable space for RAID 5 configurations is significantly higher than for the RAID 10 configurations. Usable space for RAID 5 increases further when increasing the number of disks used in the RAID layout. Hence for customers whose database performance is not paramount, opting for RAID 5 deployments is a valid tradeoff. Although RAID 10 often outperforms RAID 5 in read-and-write workloads, RAID 5 may match the RAID 10 performance in cases where the application is read-intensive

Using RAID 10 Compared to RAID 5 Configurations for Oracle® Database Deployments — Tradeoffs

or with minimal write. In such cases, RAID 5 provides customers with a greater cost benefit than the more expensive RAID 10 option.

Availability Tradeoff — RAID 10 Compared to RAID 5

This section examines the availability analysis for RAID 10 and RAID 5 implementations, providing an estimate on how long the system may take to recover or rebuild the data in the event of a disk failure.

Test Configurations

Two Oracle 11g R1 (11.1.0.7) two-node RAC clusters were deployed. One RAC cluster was deployed with a RAID 10 layout and the other cluster was deployed with a RAID 5 layout. The hardware and software components for the test were identical to the performance tradeoff test.

The difference was the way the disks were laid out for the database. Eight physical disks were used for the RAID 10 configuration and five physical spindles were used for the RAID 5 configuration. We were able to create the same amount of usable space with eight disks for RAID 10 configurations and five disks for RAID 5 configurations. The intent was to load same amount of data for both configurations and to monitor the amount of time it takes for the system to recover from a disk failure.

Test Cases

To determine how the system behaves and recovers in case of a disk failure, both databases were loaded with 400 GB of data. Then the TPC-C transaction was started for both databases using Quest Benchmark Factory. While the transactions were in progress, one disk was removed from each configuration and the data rebuild time was monitored for both the configurations.

Results and Analysis

The Navisphere GUI provides real time monitoring of the data rebuild process. After removing one disk drive, the data rebuilding was monitored for both the RAID 10 configuration and the RAID 5 configuration. The following table shows the percentage of longer rebuild time for RAID 5 compared to RAID 10, taken by the array during data rebuild process. It needs to be noted that these tests were performed with the Rebuild Priority set to **HIGH**.

Rebuild (%)	% Longer Rebuild Time for RAID 5 Compared to RAID 10
0	0%
10	35%
20	37%
30	25%
40	20%
50	20%
60	18%

Using RAID 10 Compared to RAID 5 Configurations for Oracle® Database Deployments — Tradeoffs

70	18%
80	20%
90	20%
100	24%

The RAID 10 configuration can recover faster than the RAID 5 configuration. For the RAID 10 configuration, the data was spread across all eight disks, and then the data set was mirrored. For the RAID 5 configuration, the data was spread across all five disks with the parity. The above data rebuilding trend is expected because the parity calculation for data rebuild with RAID 5 takes more time. From the table above, we can see that data rebuilding with RAID 10 may be 18% to 37% faster compared to that of the RAID 5 configuration in the event of a single disk failure.

Both of these data rebuilding tasks are non-destructive in nature. The TPC-C runs continued to run during this rebuilding process, but the performance impact was not analyzed during this data rebuilding process.

One important availability difference between the RAID 10 and RAID 5 configurations is that the RAID 10 configuration can survive more than one disk failure, whereas RAID 5 can tolerate only a single disk failure. With a RAID 10 configuration with eight physical disks, the system can survive with as many as four disk failures (without any data loss), whereas with a RAID 5 configuration with eight disks, the system can survive with only one disk failure, in other words, RAID 5 cannot rebuild the data with parity when more than one disk fails.

Four Quadrant Analysis

The four quadrant analysis can be used to determine the recommended RAID for different types of workloads and database sizes. This analysis is particularly useful for databases with mixed-use deployment, where customers run OLTP during the day and DSS workload at night. In the below chart, the X-axis represents the database size, and the Y-axis represents the type of workload. The database size can vary from small to large. The workload can vary from 100% OLTP to 100% DSS. The combination of database size and workload will decide into which quadrant the customer requirement falls. For example, in the case of a large database size with 100% OLTP workload, the requirement falls into Q1. Based on our findings, Figure 8 provides the recommended RAID types for different quadrants.

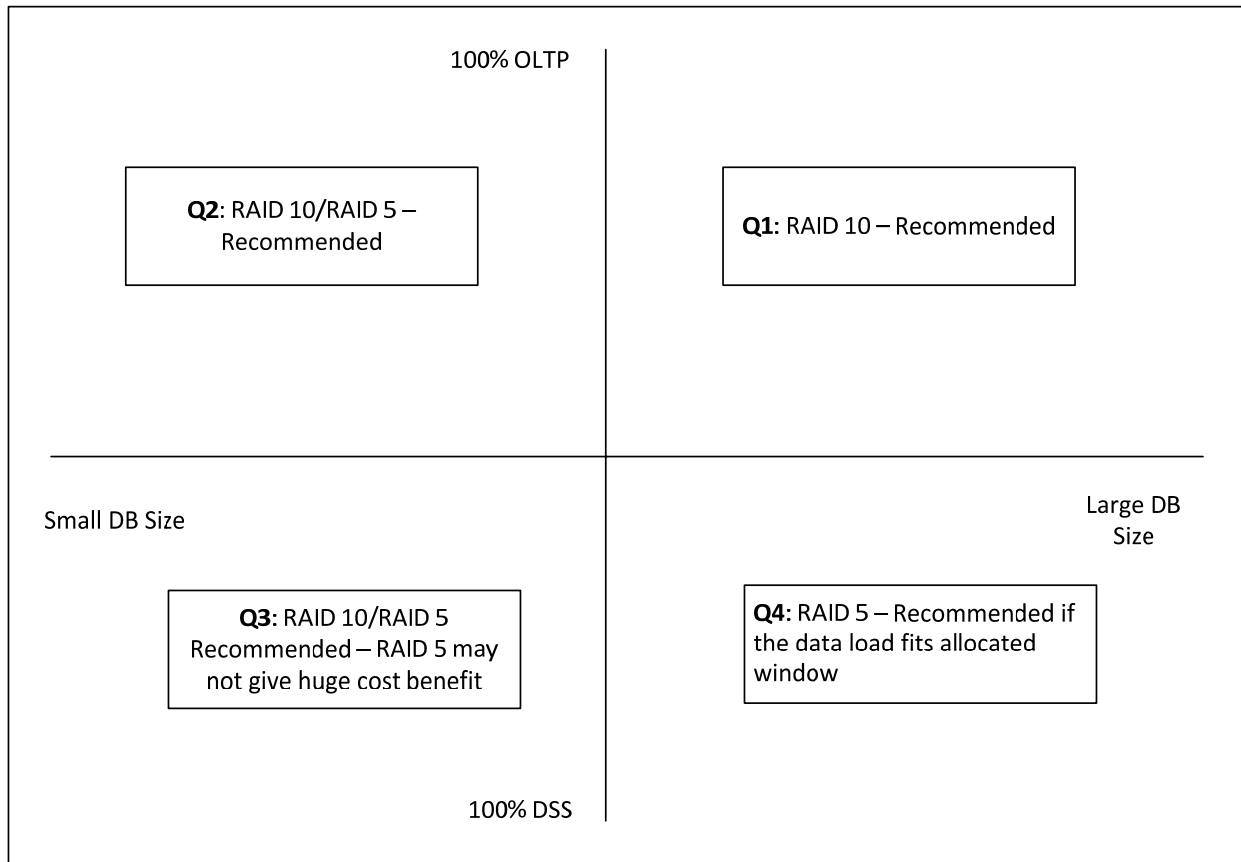


Figure 8: Four Quadrant Analysis — RAID Recommendation for Mixed-Use Database Deployments

Quadrant 1: Large Database Size and OLTP Workload

Performance and non-degrade uptime are the key considerations in this quadrant. As we discussed earlier, RAID 10 is always the best option for this quadrant.

Quadrant 2: Small Database Size and OLTP Workload

Performance and non-degraded uptime are the key considerations in this quadrant as well. Because of the small database size, the customer may not really have a huge cost benefit if choosing RAID 5. For this quadrant, although RAID 5 can work fine for the customer, RAID 10 is recommended.

Quadrant 3: Small Database Size and DSS Workload

Because of the small database size, the customer may not have a large cost benefit if choosing RAID 5. For this quadrant, although RAID 5 can work fine for the customer, RAID 10 is recommended.

Quadrant 4: Large Database Size and DSS Workload

Using RAID 10 Compared to RAID 5 Configurations for Oracle® Database Deployments — Tradeoffs

For DSS workloads, the major disadvantage is the performance bottleneck in data loading. Most of the time, the customer has a limited window to complete the data loading. If the data loading activity fits into the allocated window, RAID 5 can be best fit for this quadrant.

A RAID-type recommendation for Quadrant 4 needs further analysis, because the DSS workload can be a mix of read-and-write activities. For a mixed database usage scenario, databases are used for read-intensive queries during day time and for write-intensive data load at night. The DSS data load activity can vary from light to heavy. Figure 9 depicts the RAID-type recommendations for different read-and-write workload mixes. The X-axis represents the database size, and the Y-axis represents the workload mix.

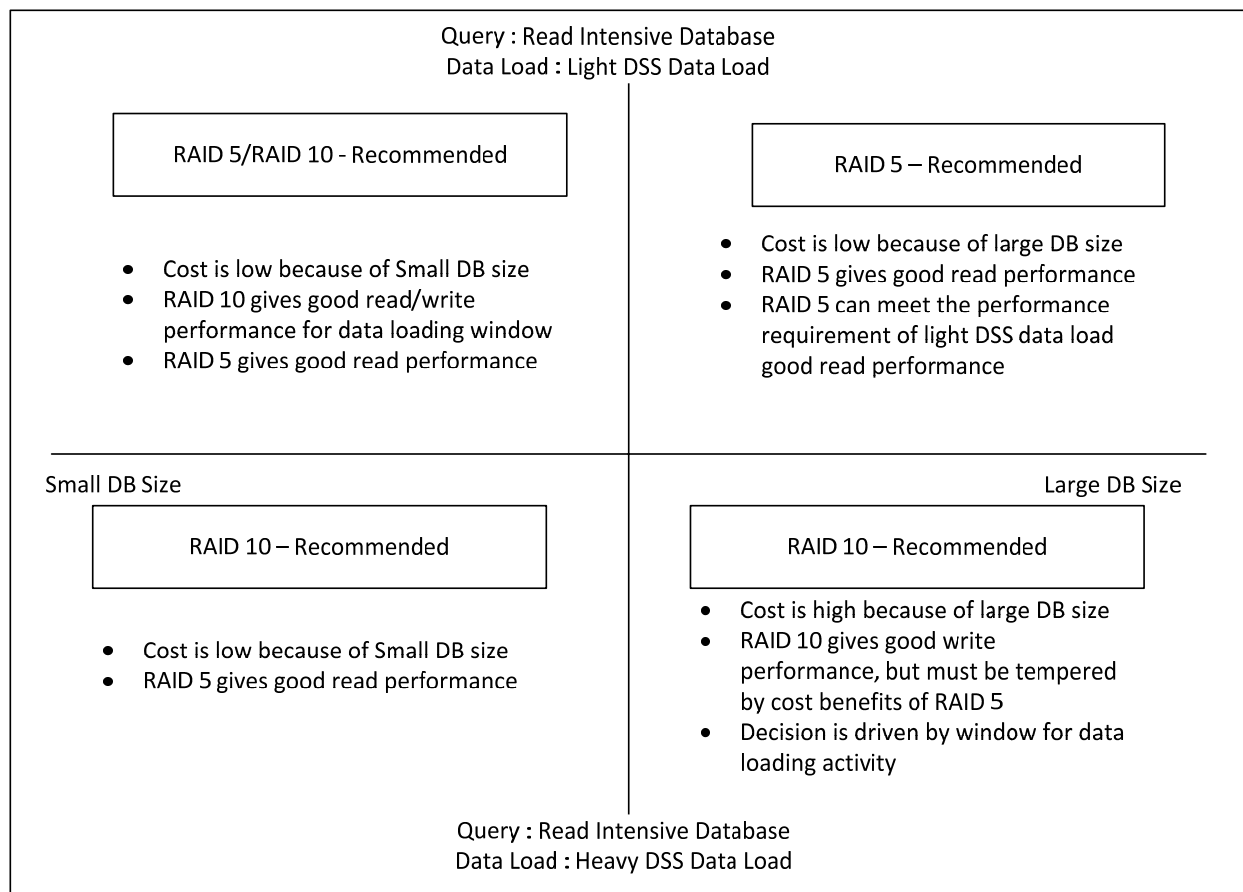


Figure 9: RAID Recommendation for Read-and-Write DSS Workload Mixes

The major consideration when designing the database for mixed use is the window for data load. The data load should be completed within the time specified, so that it does not impact the database read operations. In other words, the RAID configuration should be able to meet the DSS data-loading performance requirements. To be able to meet these performance requirements, Figure 9 can be analyzed as follows:

Using RAID 10 Compared to RAID 5 Configurations for Oracle® Database Deployments — Tradeoffs

1. For read-intensive databases with light data-load activities (Q1 and Q2 in Figure 9), RAID 5 is recommended for both small and large database sizes. RAID 5 does not create any performance bottlenecks because of the light data-loading activities. RAID 5 is a highly cost-effective solution in this case.
2. For small databases with heavy data-load activities (Q3 in Figure 9), RAID 10 is recommended. RAID 5 may be a performance bottleneck in this workload scenario. Because of the small database size, RAID 5 may not result in significant cost savings.
3. For large databases with heavy data-load activities (Q4 in Figure 9), RAID 10 is recommended. For very large databases, the RAID 10 solution will require significantly more storage and will be an expensive solution. In that case, the decision for which RAID type to select should be driven by the amount of window available for data loading.

Conclusion

Considering all analyses from the previous sections, it can be concluded that RAID 5 imposes performance penalties for write-intensive workloads. The performance penalty increases as the number of writes increases. This parity overhead may cause contention for the read operation, and this cascading effect may decrease the read performance. Hence for all write-intensive workloads, RAID 10 is always recommended to achieve optimum performance.

For read-intensive workloads, RAID 5 performance can match that of RAID 10. As RAID 5 configurations provide low cost per GB, it is recommended to have the RAID 5 layout for read-intensive workloads.

However, most of the Oracle databases activities are write intensive. The following database operations typically experience high write activities –

- Writing to Oracle Control Files
- Writing to Online redo logs
- Writing to temp tablespaces
- Writing to Undo tablespaces

Performance requirements for the above activities are very high from a database perspective. Hence from a database engine point of view, it is highly recommended to place the Oracle control files, online redo log files, temp, and undo table spaces in a RAID 10 layout.

Database data files are the target for the use of the real application. Depending on the application workload, the customer should decide to place the data files in a RAID 10 or a RAID 5 layout.

Based on the results of the testing, for OLTP workloads with a higher write to read ratio, it is recommended to select a RAID 10 layout. However, a customer can opt for a RAID 5 deployment to have a cost-effective solution. In addition, RAID 5 can be used for a lightly loaded database without much performance impact.

Using RAID 10 Compared to RAID 5 Configurations for Oracle® Database Deployments — Tradeoffs

But for write-intensive OLAP workloads (for example, database batch loading), RAID 10 is strongly recommended for better performance. For data warehouse DSS applications with highly read-intensive workloads, RAID 5 can provide very good performance as RAID 10 and, at the same time, can be very cost effective. The above recommendations are consolidated in the below table:

Oracle Database Components	Workload Type	Recommended RAID Configuration	Comments
Control files	Write-intensive	RAID 10	RAID 5 is not recommended.
Redo log files	Write-intensive	RAID 10	RAID 5 is not recommended.
Undo table space	Write-intensive	RAID 10	RAID 5 is not recommended.
Temp table space	Write-intensive	RAID 10	RAID 5 is not recommended.
Data files	OLTP (higher read-to-write ratio)	RAID 10/RAID 5	With an increasing write percentage, RAID 5 performance decreases.
Data files	DSS — write-intensive	RAID 10	RAID 5 is not recommended.
Data files	DSS — read-intensive	RAID 10/RAID 5	RAID 5 can match RAID 10 performance.

To conclude, RAID 10 delivers an optimal performance for most database workloads. However, RAID 5 performance can match the RAID 10 performance for read-intensive workloads or lightly loaded database workloads. The implementation of RAID 5 always delivers lower cost per GB, resulting in a more cost-effective solution. Overall, RAID 5 can perform as well as RAID 10 for read-intensive workloads.

Apart from the performance and cost analysis, the customer should also consider the availability perspective of the storage layout. As discussed earlier, RAID 10 can sustain more disk failures (half of the amount of disks used for the RAID configuration) with no data loss, whereas RAID 5 can sustain only one disk failure. It should also be kept in mind that the higher the amount of disks in a RAID 5 group, the more time is needed to rebuild the data in the case of disk failure.

GLOSSARY

RAID — Redundant Array of Independent Disk. The technology is used to improve performance and availability of the underlying storage arrays for any application.

OLTP — Online Transaction Processing systems which are characterized by a short response time, small transactions, and high concurrency. They usually have high performance requirements.

DSS — Decision Support Systems support business and organizational decision-making activities, and usually have high throughput requirements.

TPS — Transactions per Second indicates the transaction rate that a system can support.

Using RAID 10 Compared to RAID 5 Configurations for Oracle® Database Deployments — Tradeoffs

Response Time — Indicates the time elapsed for the SQL round trip.

Throughput — Measures the traffic flow between the storage device and the application servers.

References

Oracle® Database Performance Tuning Guide 11g Release 1 (11.1) Part Number B28274-02

http://download.oracle.com/docs/cd/B28359_01/server.111/b28274/toc.htm

Dell Enterprise Storages

<http://www.dell.com/content/products/category.aspx/storage?c=us&cs=555&l=en&s=biz&~ck=mn>

Dell PowerEdge Servers

<http://www.dell.com/content/products/category.aspx/servers?c=us&cs=555&l=en&s=biz&~ck=mn>

Dell/EMC Products

<http://www.dell.com/emc>

EMC Technical Documentations

<https://powerlink.emc.com/nsepn/webapps/btg548664833igtcuup4826/km/appmanager/km/secureDesktop?nfpb=true&pageLabel=image7a&internalId=0b014066800219ff&irrt=true>