

Evaluating Storage Performance for Oracle Database Workloads

By Russ Fellows

Sr. Analyst Performance Practice Lead

August 1, 2018



Evaluator Group

Enabling you to make the best technology decisions



Evaluator Group



NetApp

Executive Summary

Storage is a critical aspect of nearly all enterprise applications and particularly databases. Storage performance has a direct correlation to database application performance and can be one of the most efficient methods for improving database performance. The number of processors dedicated to database applications is one aspect of performance along with the amount of memory available and storage performance.

With database vendors now charging for the number of CPU cores and with available memory limited, one of the most cost-effective ways to improve database performance is by improving storage performance. The key attributes to consider for storage are response time and how cost effectively a storage system is able to provide performance and capacity. Oracle is known as a powerful enterprise database system; however, costs can quickly mount due to the per processor licensing costs.

Although IT staff understand that storage can improve database performance the difficulty often comes in how to measure storage performance in general and in particular for database workloads. One of the best methods for comparing performance is through the use of benchmarks, ideally using benchmarks that utilize workloads similar to a production database application.

The Storage Performance Council has developed two well regarded storage benchmarks, SPC-1 and SPC-2 that provide performance metrics along with cost and hence efficiency values. The SPC-1 benchmark provides information for storage when running a transactional workload, similar to many database applications. The SPC-2 benchmark ascertains performance and efficiency for data intensive applications similar to data analytics and data warehouse applications.

NetApp has been a proponent of publishing benchmarks and recently published results for the EF570^{1,2} that showed industry leading efficiency levels, as measured by price / performance ratio for both the SPC-1 and SPC-2 benchmarks. This paper examines the critical storage requirements for database applications and Oracle in particular along with aspects of storage performance that are critical to improving Oracle database performance.

Evaluator Group found that the NetApp EF570's performance and efficiency make it well suited for database workloads and in particular, the industry leading storage latency and price / performance levels make the EF570 one of the most cost-effective ways to improve Oracle database application performance. For any database that is limited by existing storage performance, the NetApp EF570 is one of the best options available, due to its industry leading performance and price / performance results.

¹ [SPC-1 Results for NetApp EF570](#)

² [SPC-2 Results for NetApp EF570](#)

Storage Requirements for Oracle

Storage has always been a key aspect of database performance and is directly linked to the availability of applications utilizing Oracle databases. DBAs are often the most demanding consumers of storage and other IT infrastructure due to their close relationship to many applications that drive modern businesses.

Many IT admins and DBAs have found that improving storage performance is the most cost-effective way to improve application performance. In addition, the efficiency benefits of high-performance storage provide compelling economic gains, often enabling companies to lower overall costs by reducing database processor licensing fees. This is a concept that can get support from CFOs along with IT administrators and DBAs.

Database licensing costs are often a significant component in total IT budget, maximizing system performance in order to reduce licensing costs is an important factor of IT optimization. As system processors and CPU cores have increased, database vendors have moved to licensing per CPU core, thereby negating the typical price / performance improvements other applications have achieved with increasing CPU cores. However, storage remains one component of systems that database applications do not charge based upon performance.

In many instances, companies utilize multiple Oracle database applications, leveraging skills, training and licensing economies of scale. Thus, for many companies the question is less about meeting the needs of one particular Oracle database application, but rather ensuring that their infrastructure choices provide high efficiency for smaller database configurations yet are also able to scale capacity and performance to meet the needs of large database applications while meeting availability requirements.

Oracle Storage Best Practices

There are a number of best practices for storage with regard to databases that should be considered along with several data protection tools commonly used. These considerations are summarized below:

- **Storage for Oracle Databases**
 - **ASM** – Provides DBA directed control over storage resources
 - **Storage Performance** – Latency is a key aspect, along with high I/O rates for random I/O
 - **SAME** – Stripe And Mirror Everything is an adage that still applies, albeit with considerations for the capabilities of high-performance external storage
- **Oracle Data Protection Tools**
 - **RMAN** – Is the basis for application consistent Oracle backup and recovery, and can utilize high-performance SAN storage as targets, with incremental forever and data compression
 - **Fast Recovery Area (FRA)** – Recommended for maintaining and managing disk-based backups, uses change block tracking for incremental backup tiering to storage
 - **Flashback** – Leverages FRA to log change tracking in order to restore a database to a point in time, thereby recovering from user errors, corruption or other logical errors
 - **Active Data Guard (ADG)** – Application Continuity using Replication with RMAN block change tracking for management, monitoring and automation of database protection

The storage features that are most critical in an Oracle environment are reliability, performance and scalability. While other storage features may provide benefit, performance and reliability underlie all of Oracle's application level availability and data protection features. As a result, many DBAs prefer to control their data protection and DR strategies utilizing Oracle tools.

The listed recommendations for Oracle have remained consistent and many of the data protection features provide a trusted set of tools that enable database administrators to deliver consistent application and data protection that is hardware independent. With a variety of Oracle tools to protect from physical media errors, logical errors or disasters, additional data protection mechanisms are often unnecessary or create administrative overhead.

Additional Storage Considerations for Oracle

With the technologies previously listed, particularly RMAN, FRA, Flashback and ADG many of the data protection tasks for Oracle can be managed via these tools. Moreover, although there is a role for disk-based snapshots and replication, in many instances these capabilities are used to supplement Oracle tools. Other features, such as data reduction storage capabilities may provide little benefit. The storage features which should be considered optional for Oracle include the following:

- **Storage replication** – May be utilized to supplement Oracle Active Data Guard, but may not be the first choice for DBAs
- **Storage data reduction** – Has little value because Oracle data does not deduplicate and Oracle can manage table compression both in tables and for backup
- **Storage snapshots** – May be beneficial for creating point in time copies; however, RMAN or FRA with Flashback guarantees application consistent point in time copies

NetApp EF570

The NetApp EF570 is an all-flash storage system that contains many of the qualities that make it an ideal fit for use with a variety of applications and in particular a good fit for database applications. An overview of some of the important storage system requirements for Oracle database applications is provided below in Table 1 highlighting the requirement, a relevant EF570 feature and what benefit that provides to system and database administrators.

Shown below in Table is an overview of Oracle Storage requirements along with NetApp EF570 features and their benefits for Oracle environments.

Oracle Storage Requirement	NetApp EF Feature	Oracle DBA Benefit
Reliability	Fourth generation system with documented six nines of availability	Proven product and support by a leading vendor provides a trusted storage choice
Performance	All flash EF570 delivers sub 300 us latency, 1M 4K IOPS and 16 GB/s bandwidth with only 24 SSD's	High performance entry configuration, additional headroom with expansion SSDs, scale-out performance with Oracle ASM
Scalability	Scale to 120 SSDs and 1.8 PB online, plus ability to manage multiple EF systems from one management console	No disruptions when scaling storage capacity up and out to meet growing capacity or performance needs
Configuration Options	Online Volume Groups, RAID, Cache and other storage configuration options	Ability to adjust storage performance without disruption or additional capacity requirements
Data Integrity	Mirrored controller write cache for 100% integrity with T10-PI data assurance	HA design with end-to-end integrity checks between Oracle and media ensures no data loss or corruption
Database Availability	Synchronous and Asynchronous Replication with multiple consistency groups	Flexibility to meet different database recovery point objectives between local and remote sites
Storage Protection Options	EF Snapshots, with Sync and Asynch replication enable using all-flash for primary and mixed-media E-Series systems for data protection	Array based protection options can supplement Oracle tools and utilize SSD's or high-capacity HDD's for snapshot or replication targets when desired

Table 1: NetApp EF570 Features & Benefits for Oracle Workloads

NetApp EF570 Performance

Claims of high performance abound in marketing literature and for storage products this is particularly true. The challenge for DBA's, architects and IT administrators is to understand what performance claims are based on measurable results and facts, rather than marketing conjecture. In some instances, performance claims may be based upon real data, but do not utilize methods that allow results to be compared to other systems. In other instances, vendors make no specific claims, instead opting to provide a "guarantee" that performance will be good enough. Typically, vendors are not intentionally withholding information, they just have not invested the resources required to publish benchmark results. In other cases, vendors are attempting to disguise unfavorable data.

Thus, whenever a vendor publishes relevant benchmark results, this should be seen as a significant sign of confidence and an investment the vendor has made by publishing data that can be helpful to consumers. NetApp has been one of the most prolific vendors in publishing benchmark results for their systems since the inception of the company. NetApp recently published results for the EF570 which achieved the highest price / performance results on both the Storage Performance Council's SPC-1 and SPC-2 benchmarks when released in September 2017.

The SPC-1 benchmark has published hundreds of system results since its inception. Although the SPC-1 workload does not utilize an actual database, the workload is very similar to many databases and was designed to be representative of a broad class of transactional applications. Moreover, SPC-1 is a highly relevant benchmark for comparing results of storage systems when running a transactional, database like workload. The SPC-2 benchmark is designed to test a storage systems ability to perform large data transfers, similar to those found in data warehouse and data analytics, technical computing and video capture.

EF570 Benchmark Results

The NetApp EF570 achieved high performance and one of the best price / performance ratios of any system reporting SPC-1 results while also having the lowest overall response time. These factors are critical for improving Oracle storage performance. The two most critical factors are the response time and performance efficiency, with the cost for usable capacity another important factor. The overall performance level is a consideration; however, when utilizing Oracle ASM total storage performance levels can easily be scaled up as desired. As discussed previously, response times and performance efficiency do not improve by scaling storage, which is why these factors are the most critical storage performance evaluation criteria for transactional workloads.

NetApp's SPC-2 results for the EF570 highlight the systems performance capabilities when running data intensive applications that depend on high data transfer rates. By establishing leading price / performance results for both benchmarks, the EF570 is a strong contender for applications that demand the lowest latency and for those that demand high data rates.

EF570 SPC Benchmark Highlights:

- The EF570 attained the best performance efficiency results for both the SPC-1 and SPC-2 workloads
- The EF570 had the best (i.e. lowest) response time (measured by ms of latency) for the SPC-1 benchmark
- Specifically, EF570 system achieved the following SPC-1 results:
 - Price / Performance = \$128.42 per 1,000 SPC-1 IOPS
 - Total performance of 500,022 SPC-1 IOPS, at a maximum 0.437 ms latency
 - Total configuration cost was \$64,212.58
 - Cost per usable GB was \$7.13, utilizing 24 SSDs at 800 GB with 19.2 TB of raw capacity
- The EF570 achieved the following SPC-2 results:
 - Price / Performance of \$3.69 / SPC-2 MBPS
 - An aggregate score of 17,337 SPC-2 MBPS,
 - A cost of \$63,924.52 and a usable capacity of 12.708 GB

Oracle Database I/O

For transactional Oracle database applications, storage must provide very low latency random reads in 8KB and larger I/O's while also supporting large block sequential write operations fast enough to not cause any queuing or delays. The recommended storage layout is to utilize a minimum of two devices for the database and a separate device for storing logs, although more volumes are recommended for larger databases or when higher performance is required.

Oracle DBA's typically utilize ASM as the underlying mechanism for managing database storage utilization, in effect acting as a logical volume manager which eliminates the need for 3rd party volume managers or filesystems. However, ASM is not directly in the data path, as all I/O is performed by the database to the raw block mapping managed by ASM. Thus, traditional Oracle I/O mechanisms remain. The default block size setting in Oracle is 8 KB, which represents the smallest I/O that may occur. Typically, data is accessed in stripes, which range from multiple 8 KB blocks up to 64 MB stripes when ASM is load balancing or mass data updates. Moreover, storage systems must provide high performance for blocks as small as 8KB and up to a megabyte or larger. In particular, random read performance is critical, followed by random and sequential write speed.

For data analytic workloads, including Oracle warehouse applications, I/O is performed significantly differently, utilizing larger I/O transfer sizes and with a nearly sequential access pattern. Additionally, reads and writes are typically not interspersed, enabling the storage system to optimize data transfer operations. Rather than latency being the primary consideration, the data transfer rate is the primary consideration.

SPC Benchmark Details

Benchmarks are an effort to provide verifiable information that enable IT consumers to make educated decisions. The Storage Performance Council was created over a decade ago to help address the dearth of publicly available storage benchmarks, providing relevant data with verified performance and pricing. The SPC-1 benchmark was designed to measure storage performance for applications that were transactional

in nature. Although SPC-1 does not utilize a database, the benchmarks design generates I/O access patterns that are similar to the way many databases access storage.

The SPC-1 benchmark utilizes multiple random, I/O streams to the two primary volumes, and a single, write only stream to the third volume with specific I/O streams for each application storage unit or ASU. The block size of I/O requests varies from 8 KB up to 128 KB, with 8 KB comprising some 80% of all I/O requests. The first ASU is referred to as a “Data Store”, the second as a “User Store” and the third as the “Logging Write Store.” The block size ratio, and the number of read and write operations is distinct for each storage region or ASU.

Overall, the SPC-1 benchmark produces a roughly 50% read to write ratio, although this varies per region and includes hot spots and other access patterns meant to simulate accessing an index and other database patterns. Clearly the designers of the benchmark are attempting to recreate a storage layout and workload that are similar to that found in typical database workloads.

A small storage configuration for an Oracle database typically consists of two storage volumes for the primary database and a third volume for redo logs. Separating the sequential write access of the redo logs from the main database eliminates contention which improves transactional performance and logging. The SPC-1 benchmark utilizes this same storage layout, requiring two application storage units (ASU’s) optimized for low latency, which together comprise 90% of total capacity and a third smaller ASU of 10% total capacity that is optimized for sequential writes.

Final Thoughts

It is important to first understand the nature of a particular workload or set of database workloads to determine the point of contention, or bottleneck. For databases that are CPU bound, almost no amount of storage increase will alleviate the bottleneck. Memory contention can be improved by some degree with faster storage, since storage is a substitute for memory in most database designs, although it is significantly slower. However, for database applications that are limited by I/O, moving to a high-performing, low-latency storage system can provide significant performance improvements.

The SPC-1 benchmark provides a highly relevant data point for estimating Oracle transactional database performance. Although it is difficult to provide absolute values, it is appropriate to use SPC-1 results as a relative performance factor. A system that achieves 2X better performance, or price / performance results on SPC-1 than a competitor would be expected to have similar benefits for an Oracle workload that was I/O bound. Another, related by separate factor is the latency. If two storage systems can produce the same I/O rates, but one does so at 2X lower latency, that will also factor into improving overall database performance for any I/O bound workload.

The NetApp EF570 is one of the few systems to publish both SPC-1 and SPC-2 workloads for the same system. This is unusual because the two workloads place very different demands on the system. Typically, servers and storage are optimized for a particular type of workload, since it is rare that a system can perform both efficiently. The all-flash NetApp EF570 is able to perform a variety of workloads at world

record levels, with the only change being the RAID level utilized. The EF570 is a proven mid-range system with a strong pedigree of delivering outstanding performance cost effectively.

For organizations running smaller Oracle database applications, the NetApp EF570 is likely to support multiple applications, delivering consistent performance, cost effectively. For larger organizations running many Oracle database instances or large configurations including Oracle RAC, multiple EF570's may be utilized to deliver scale-out storage performance and capacity together with Oracle ASM.

About Evaluator Group

Evaluator Group Inc. is dedicated to helping [IT professionals](#) and vendors create and implement strategies that make the most of the value of their storage and digital information. Evaluator Group services deliver [in-depth, unbiased analysis](#) on storage architectures, infrastructures and management for IT professionals. Since 1997, Evaluator Group has provided services for thousands of end users and vendor professionals through product and market evaluations, competitive analysis and [education](#). www.evaluatorgroup.com Follow us on Twitter [@evaluator_group](#)

Copyright 2018 Evaluator Group, Inc. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written consent of Evaluator Group Inc. The information contained in this document is subject to change without notice. Evaluator Group assumes no responsibility for errors or omissions. Evaluator Group makes no expressed or implied warranties in this document relating to the use or operation of the products described herein. In no event shall Evaluator Group be liable for any indirect, special, inconsequential or incidental damages arising out of or associated with any aspect of this publication, even if advised of the possibility of such damages. The Evaluator Series is a trademark of Evaluator Group, Inc. All other trademarks are the property of their respective companies.