



Technical Report

# **Oracle Database Performance on NetApp ASA**

Authors – Bobby Oommen, Joe Carter - NetApp

July 2025 | TR-5015

## **Abstract**

This technical report details the NetApp ASA performance for an Oracle database which was configured in the most optimal way to ensure the maximum throughout the solution validation.

## TABLE OF CONTENTS

<b>Introduction .....</b>	<b>3</b>
<b>Executive Summary.....</b>	<b>3</b>
<b>ONTAP ASA platforms .....</b>	<b>3</b>
<b>ASA performance .....</b>	<b>3</b>
FCP 100% Select.....	4
FCP 75% Select, 25% Update.....	5
Database servers.....	6
Network.....	6
Oracle ASM configuration .....	6
Oracle SLOB.....	7
Working sets .....	7
Curve generation .....	7
Data selection .....	7
<b>Where to find additional information.....</b>	<b>8</b>
<b>Version History .....</b>	<b>8</b>

# Introduction

NetApp® ONTAP® is a powerful data-management platform with native capabilities that include inline compression, nondisruptive hardware upgrades, and the ability to import a LUN from a foreign storage array. Up to 24 nodes can be clustered together, simultaneously serving data through Network File System (NFS), Server Message Block (SMB), iSCSI, Fibre Channel (FC), and Nonvolatile Memory Express (NVMe) protocols. In addition, NetApp Snapshot technology is the basis for creating tens of thousands of online backups and fully operational database clones.

In addition to the rich feature set of ONTAP, there are a wide variety of user requirements, including database size, performance requirements, and data protection needs. Known deployments of NetApp storage include everything from a virtualized environment of more than 6,000 databases running under VMware ESX to a single-instance data warehouse currently 996TB in size and growing.

The focus of this document is bare metal database performance using NetApp ASA storage systems.

## Executive Summary

This document serves as a guide for IT professionals, architects, and decision-makers looking to understand and enhance storage and database performance in environments leveraging ONTAP ASA platforms. The Technical report details in-depth exploration of storage performance on ONTAP ASA platforms, with a particular emphasis on optimizing database environments examining the KPIs relevant to modern database workloads offering insight into how latency affects application responsiveness and user experience. Also, the document offers best practices and analytical perspectives aimed at maximizing storage performance for demanding mission critical enterprise database applications.

## ONTAP ASA platforms

ONTAP software is the foundation for advanced data protection and management. However, ONTAP doesn't only refer to software. There are multiple ONTAP hardware platforms to choose from that rely on a variety of storage technologies including flash media, spinning drives, and virtualized storage. Nearly all databases deployed today are hosted on solid-state storage, and the trend is accelerating.

NetApp ASA provides a robust, scalable, and secure infrastructure foundation designed to support enterprise application workloads with high performance, availability, and centralized management. The ASA framework leverages modular design principles, enabling seamless integration of compute, storage, and network components within a validated, interoperable ecosystem. When deployed with SAN (Storage Area Network), ASA decouples storage from compute, allowing shared, block-level access to storage resources via Fibre Channel or iSCSI, optimized for reliability and speed.

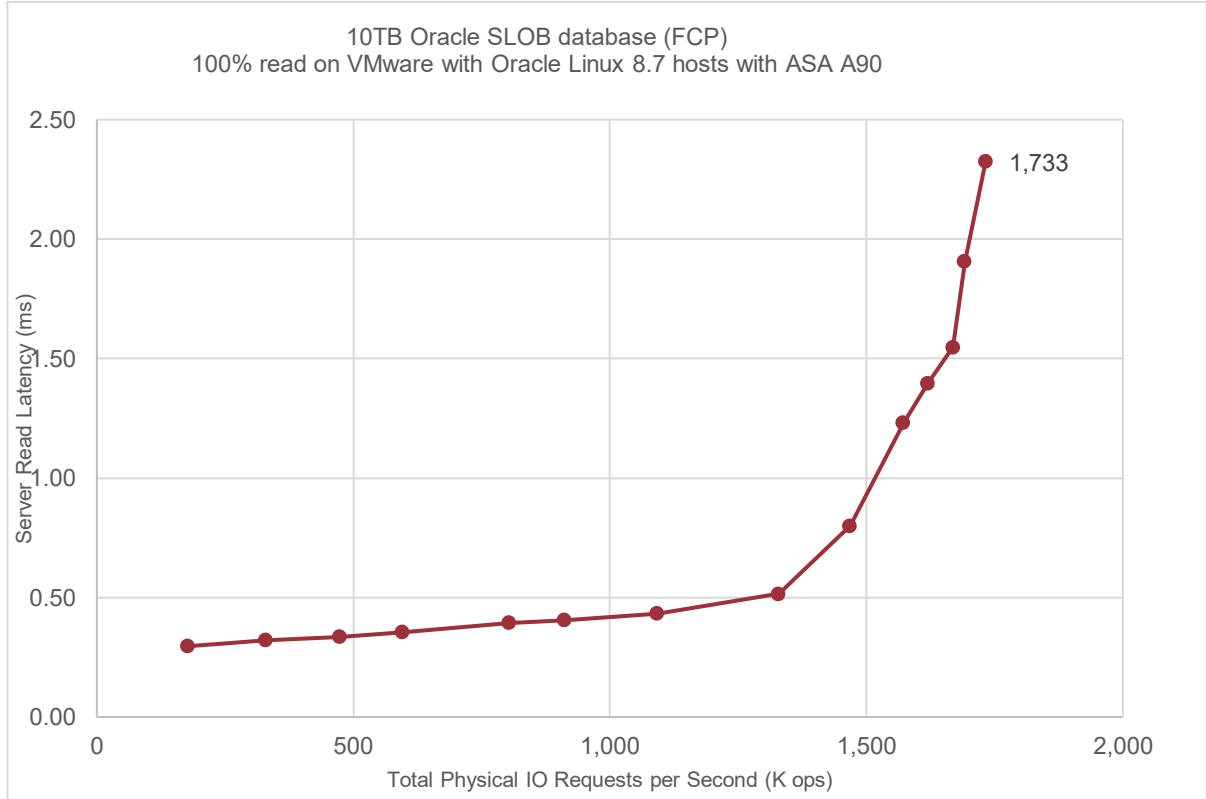
The architecture supports secure multi-tenancy through ONTAP Storage Virtual Machines (SVMs), ensures consistent data protection via technologies like Snap Center, and integrates easily with leading virtualization platforms such as VMware. Best practices around SAN zoning, LUN masking, multipathing, and consistency groups are enforced to guarantee operational resilience and data integrity. Additionally, compatibility with industry-standard hardware and adherence to NetApp's Hardware and Software Universe (HWU) ensures supportability and lifecycle alignment. Overall, NetApp ASA delivers an efficient, resilient, enterprise-grade solution tailored for mission-critical applications and modern data center environments.

## ASA performance

The following graphs illustrate the performance capabilities of the NetApp ASA platforms for several configurations. The tests utilized Oracle 21c standalone databases with ASM to generate the workload and were conducted on representative two-node ASA A90 HA pairs.

## FCP Validation


### FCP 100% Select



#### The highlights of 100% selects are as follows:

The graphs show the performance capabilities up to the point where the configuration saturates. The ASA controllers maintain latencies and start times under 300 $\mu$ s, remaining below 500 $\mu$ s as IOPS approaches the 1M mark with storage efficiency, 100% data availability, and the ability to ensure 100% data integrity. These specifications adequately meet the requirements for any mission-critical applications. When users were scaled to saturate the system, an IOPS of 1.73 million with a latency of 2.3ms was recorded. Consequently, NetApp ASA controllers frequently eliminate storage as a database performance bottleneck, leaving limitations primarily to query logic and CPU processing on the server while providing the highest availability and data protection.

## FCP 75% Select, 25% Update

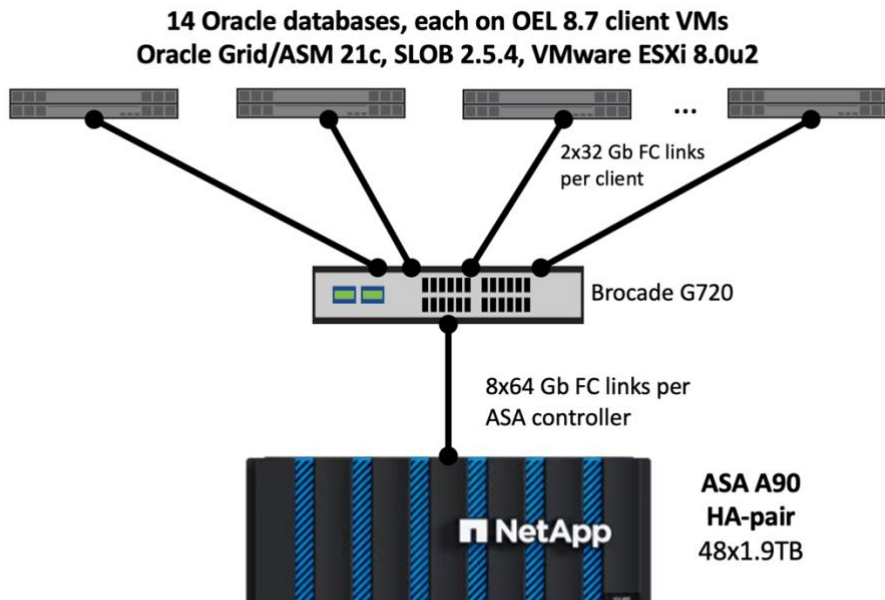


### The highlights of 75% selects & 25% updates are as follows:

The graph above shows a 75% SELECT test, which results in about an 80% read IO ratio. The reason is a test involving 25% update operations creates a read for each block updated, which slightly increases the read percentage. The latencies are a little higher than the 100% read test, which is to be expected as the write I/O is producing additional load on the controller. The storage OS is still ONTAP, which means the same write path is used when changes are committed to NVRAM, staged in RAM, and then written out to the drives.

There are many ways to measure write latency at the database level. Block sizes, parallelism, and the system calls used to perform the write vary. NetApp generally focuses on `log file parallel write` latency. This is the storage I/O component of a database redo logging operation, which is usually the most latency-sensitive aspect of database performance. DBAs sometimes focus on `log file sync` latency, but this is not entirely a storage operation. It will trigger a `log file parallel write event`, but it also includes other database-level operations that can delay the eventual storage I/O component of a redo log commit.

## Test configuration



A separate SLOB database was created on each of 14 databases for a total of 10TB of SLOB data.  
All hosts zoned to all LUNs/namespaces with multiple paths.  
One VMDK datastore is created per LUN.

## Database servers

Oracle 21c Grid Infrastructure and Oracle 21c Database were used for these tests.

One virtual machine was created on each of 14 ESXi 8.0u2 servers. Each VM has Oracle Enterprise Linux 8.7. Physical hardware for each ESXi node was Fujitsu PRIMERGY RX2540 M4, 256GB, Emulex Lpe32002-M2. One SLOB database was created on each of the 14 VMs. The total amount of SLOB data in the working set for this test was 10TB.

Oracle datafiles used by SLOB were initialized using the OBFUSCATE feature. This option was selected to intentionally randomize data and hence reduce the storage controller's ability to compress and deduplicate data.

## Network

The FC network was configured with two 32Gb FC connections for each database server and eight 64Gb FC connections for each storage controller node. As controllers are normally deployed in HA pair, we used a total of sixteen FC connections for the ASA A90 HA-pair.

The IP network used 10Gb NICs. The workload generator requires minimal server-to-server communication.

## Oracle ASM configuration

With 14 database servers in this test, we wanted to ensure we had enough LUN containers for ASM disk groups, but also keeping a multiple of 14.

Each database contained 3 ASM disk groups: Data, Logs and Infra.

The SLOB tablespace was on the Data disk group. This disk group contained 8 LUNs per server, or 112 total in the test bed.

For the Logs diskgroup, we selected 2 containers per database. This allowed each database server to have one logs container on each storage controller.

A VMware datastore was created for each raw device which is mapped from NetApp storage LUN, and was used to create the disks making up the ASM disk groups on each server.

## Oracle SLOB

SLOB ([Silly Little Oracle Benchmark](#)) is the premier tool for I/O benchmarking with an Oracle database. Other tools such as HammerDB or Swingbench are either too complicated to set up or have dependencies on the server hardware, Oracle version, or Oracle configuration. SLOB is the ideal tool to drive a database to perform storage I/O, and the results are a product of the complete storage path, meaning the storage-related parameters on the database, the network characteristics and limits, and of course the storage IOPS and response time capabilities.

All tests were run with a very small Oracle buffer cache, which ensures minimal cache hits on server RAM and maximizes storage IO on the array.

Finally, the SLOB tables were populated using `OBFUSCATE_COLUMNS=TRUE`. This setting is critical to obtaining accurate results with modern storage systems such as ONTAP that include efficiency features such as compression and deduplication. Without this setting, the data created by SLOB is extremely and unrealistically compressible and deduplicatable, which can lead to excessive caching on the storage system. The goal of these tests is to stress the complete code path between the Oracle database and the drives, and not to test cached IO performance.

## Working sets

The working set refers to the part of a dataset that is accessed frequently. In many Oracle databases, random I/O is concentrated in a small area of the total database size. Benchmark tests and performance results often involve small working sets where almost all I/O is handled directly from RAM rather than backend storage drives. This can affect the accuracy of the results. To prevent this skewing in these tests, it was ensured that most of the I/O involved actual drive reads.

We also did not vary the number of SLOB users during the tests. Understanding why requires understanding the difference between a SLOB user and a SLOB thread. When you create tables for SLOB testing, you must specify the number of users (also referred to as schemas). If you create a SLOB database that is 1TB in size with 16 users, the result is effectively 16 partitions of that 1TB space, each of which is 64GB in size. Tests involving only a single user or schema would exercise only 64GB of space.

As a result, testing that increases the overall IOPS on the system by increasing the number of users or schemas also expands the working set size. The tests with fewer users or schemas have a much higher proportion of cache hits on the storage system, and thus the latency looks better than it would have in a more realistic configuration. A more realistic approach is to always use the SLOB schemas but vary the threads per schema. This is supported by SLOB.

Therefore, we created a balanced configuration. We built 10TB tables for use by SLOB, split evenly across 14 databases. We set the SLOB SCALE parameter to 187246M and used 4 schemas per database. All tests involved all 4 schemas from 14 active databases to ensure that the entire 10TB database was being accessed at every load point.

## Curve generation

The performance graphs were created by performing 10-minute tests at increasing loads until the saturation point was reached. Each test was run three times to validate that the results were consistent.

## Data selection

Results were reported in terms of IOPS and latency. A SLOB thread functions solely as a load injector, and the IOPS produced by a defined number of SLOB threads can vary based on the RAC configuration,

Oracle version, and other factors. It is not possible to extrapolate SLOB users to something applicable in practical scenarios. The IOPS and associated latencies are the primary factors of importance.

The X-axis shows total datafile reads and writes. Tests with over 0% writes involve redo logging activity, generating significant I/O that is difficult to measure as IOPS due to varying block sizes and asynchronous or synchronous I/O. We focused on the most demanding I/O types: datafile random read IOPS and datafile random overwrites.

The Y-axis in the initial graphs show the random read latency, which is usually the most important number because random read latency is the primary performance limit for a database, assuming storage I/O is a source of performance limits for a given database. We take this number from the Oracle `db file sequential read` statistic. Despite its name, this is not sequential I/O. It is a random I/O operation against an indexed sequence of blocks, sometimes defined as a sequence of block reads.

We also provided random block overwrite latency on some graphs. This was taken from the Oracle `db file parallel write` statistic. This type of I/O is almost never a limiting factor for databases because it is almost always background operations. Changes to the database are committed to the redo logs, and the datafiles are updated later. We only provided the random write latency to demonstrate that the write behavior is largely the same on the A-Series and C-Series because the write path is the same. Writes are committed to the NVRAM, not the backend disk.

In addition, the storage write latency is better than this number suggests. This type of I/O is usually performed by the database as asynchronous batches of I/O. The individual I/O operations for each block would show lower latency if that level of detail could be measured from the database level.

## Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

- NetApp ASA  
<https://www.netapp.com/all-flash-san-storage-array/>
- NetApp ASA Documentation  
<https://docs.netapp.com/us-en/asa-r2/index.html>

## Version History

As an option, use the NetApp Table style to create a Version History table. Do not add a table number or caption.

Version	Date	Document Version History
Version 1.0	July 2025	Initial release



Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

### **Copyright information**

Copyright © 2023 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications. LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data—Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

### **Trademark information**

NETAPP, the NETAPP logo, and the marks listed at <https://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.