

File Storage Service Performance Guide

ORACLE WHITE PAPER | JUNE 2019





Disclaimer

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Revision History

The following revisions have been made to this white paper since its initial publication:

Date	Revision
June 19, 2019	Made minor updates and additions
September 28, 2018	Initial publication

You can find the most recent versions of the Oracle Cloud Infrastructure white papers at <https://cloud.oracle.com/iaas/technical-resources>.



Table of Contents

Overview	4
What Is File Storage?	4
Testing Methodology	4
File Storage Read/Write Performance	4
File Storage Metadata Performance	5
Maximizing File Storage Performance	5
Conclusion	6



Overview

This paper describes the performance of the Oracle Cloud Infrastructure File Storage service. Specifically, it describes what you can expect for read, write, IOPS, and metadata performance, and it offers suggestions for optimizing performance. This information can help you assess whether File Storage is suitable for your workloads.

What Is File Storage?

Oracle Cloud Infrastructure File Storage is a fully managed file storage service that can be accessed concurrently by thousands of compute instances. Using industry-standard Network File System (NFSv3) file access protocols, you can start with a file system that contains only a few kilobytes (KB) of data and scale to exabytes (EB) of data without capacity planning or provisioning. File Storage protects your data by maintaining multiple replicas, providing transparent on-disk encryption, and enabling frequent snapshots. File Storage provides high durability and high availability for your data with the consistency model of traditional NFS filers.

Testing Methodology

To test File Storage performance, we used internally developed tools for both aggregate and individual results. We plan to expand our performance testing to include industry-standard benchmarks in the future.


File Storage Read/Write Performance

File Storage is a cloud storage service—the available read and write throughput increases proportionally to the size of the file system. For each terabyte (TB) of data stored, customers who are performing reads and writes of large (~1 MB) blocks can expect the following results:

- Overall read performance of at least 100 MB/second
- Overall write performance of at least 50 MB/second
- At least 2,500 read IOPS

Currently, we don't reserve or throttle performance; therefore, observed performance might be higher or lower for your scenarios. Burst performance might be significantly better than these targets.

The highest levels of performance assume concurrent access and can be achieved only by using multiple clients, multiple threads, and multiple mount targets. A multithreaded client running on a single bare metal instance with a single mount target should be able to achieve approximately the performance described for a 10-TB file system.



The following table describes the level of performance that customers can obtain for different size file systems. Although we don't guarantee these numbers, we expect that customers who follow best practice can achieve this level of performance in the current File Storage service.

Data Set Size	Read Bandwidth (1-MB Block)	Write Bandwidth (1-MB Block)	Read IOPS (8K Block)
1 TB	100 MB/s	50 MB/s	2,500
10 TB	1 GB/s	500 MB/s	25,000
100 TB	10 GB/s	5 GB/s	250,000

For advice on reaching this level of scalability, see the “Maximizing File Storage Performance” section.

File Storage Metadata Performance


The performance of metadata operations such as file create, delete, and rename is difficult to characterize in a single number. Although File Storage runs independent metadata operations in parallel, many commonly used applications don't issue parallel requests. Therefore, the performance of a client performing operations such as `ls` or `tar` is limited by the round-trip latency of the system. The following common examples illustrate this performance:

- Running `ls -l` (long format) for 100,000 files in one directory takes 12.7 seconds. Running `ls` (short format) takes 8.3 seconds.
- Running `tar -xf` to unpack the Linux kernel (834 MB, about 67,000 files) takes 12 minutes, 30 seconds.
- A parallel implementation of `tar -xf` can unpack the Linux kernel in 60 seconds with 32 threads.

Maximizing File Storage Performance

To optimize the performance of File Storage, consider the following guidelines:

- File Storage performance increases with parallelism. Increase concurrency by using multiple threads, multiple clients, and multiple mount targets. In particular, scalability will be greatest when clients and threads are accessing independent portions of the file system.
- Use tools to run file operations in parallel. The File Storage engineering team has developed parallel tar and untar (`puntar`), parallel copy (`parcp`), and parallel remove



(`parmm`) tools. These tools are available in the `fss-parallel-tools` package in Oracle Linux.

- The available bandwidth to a file system can significantly impact its performance. In Oracle Cloud Infrastructure, larger instances (more CPUs) are entitled to more network bandwidth. File Storage performance is best with Oracle bare metal instances or large VM shapes.
- To minimize latency, clients, mount targets, and file systems should be in the same availability domain.
- For best performance, don't set any mount options such as `rsize` or `wsize` when mounting the file system. In the absence of these options, the system automatically negotiates optimal window sizes.
- Due to the limitations of Oracle Cloud Infrastructure's VNICs, each mount target is limited to about 600 MB/s of read or write traffic. If you have bandwidth-heavy workloads, consider spreading your workload across multiple mount targets after your file system exceeds 10 TB.

Conclusion

If you need highly scalable, highly durable, and highly available storage, you can benefit from using the File Storage service. The performance of File Storage scales linearly with capacity, so as you add more data you can expect a corresponding growth in performance. The elasticity of File Storage eliminates capacity planning and growth management, significantly reducing both storage and operational costs.

The following types of workloads are ideal for the File Storage service:

- Enterprise applications that need consistent file systems with the ability for shared access by multiple clients
- Applications with parallelized workloads and for performing bulk operations
- Any application that needs scale-out storage space or scale-out access

Additionally, if you want to eliminate storage maintenance tasks such as patching, security and software upgrades, capacity planning, disk failure, and handling failover, File Storage is the right solution. If you want more guidance or assistance on performance-related questions, send an email to mona.khabazan@oracle.com.







Oracle Corporation, World Headquarters

500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries

Phone: +1.650.506.7000
Fax: +1.650.506.7200

CONNECT WITH US

-  blogs.oracle.com/oracle
-  facebook.com/oracle
-  twitter.com/oracle
-  oracle.com

Integrated Cloud Applications & Platform Services

Copyright © 2019, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0719

File Storage Service Performance Guide
June 2019
Author: Oracle Corporation

