

ORACLE

# Deep Dive into Oracle Unified Directory 12.2.1.4.0 Performance

---

January 2023, Version 1.0  
Copyright © 2023, Oracle and/or its affiliates  
Public

## Purpose statement

This document discusses performance recommendations and sizing for Oracle Unified Directory as part of Oracle Identity Suite, Release 12.2.1.4.220405.

## Disclaimer

This document in any form, software or printed matter, contains proprietary information that is the exclusive property of Oracle. Your access to and use of this confidential material is subject to the terms and conditions of your Oracle software license and service agreement, which has been executed and with which you agree to comply. This document and information contained herein may not be disclosed, copied, reproduced or distributed to anyone outside Oracle without prior written consent of Oracle. This document is not part of your license agreement nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

This document is for informational purposes only and is intended solely to assist you in planning for the implementation and upgrade of the product features described. 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 in this document remains at the sole discretion of Oracle. Due to the nature of the product architecture, it may not be possible to safely include all features described in this document without risking significant destabilization of the code.

## Table of contents

---

<b>Purpose statement</b>	<b>2</b>
<b>Disclaimer</b>	<b>2</b>
<b>INTRODUCTION</b>	<b>4</b>
<b>ORACLE UNIFIED DIRECTORY BENCHMARKS OVERVIEW</b>	<b>4</b>
Test Flow Scenarios	4
OUD Search Scenarios:	4
Combined Load Scenarios:	4
Topology	5
Hardware Configuration	5
<b>ATOMIC SEARCH BENCHMARKS</b>	<b>6</b>
<b>COMBINED SEARCH BENCHMARKS</b>	<b>8</b>
<b>PERFORMANCE TUNING RECOMMENDATIONS</b>	<b>8</b>
OUD Tuning	9
OUD JVM Tuning	10
<b>CONCLUSION</b>	<b>11</b>
<b>GLOSSARY</b>	<b>11</b>

## INTRODUCTION

With increasing requirements for high scalability and performance in the field of Oracle Unified Directory, Oracle has conducted scalability benchmarks for OUD as part of the Oracle Identity and Access Management Release 12.2.1.4.220405. The environments for OUD scalability benchmarks were deployed on virtual Compute and DBaaS shapes in Oracle Cloud Infrastructure. The objective was to measure performance in specific scenarios including exact, substring, and complex filters with substring searches as well as mixed load scenarios which mimic the majority of the usage of OUD.

## ORACLE UNIFIED DIRECTORY BENCHMARKS OVERVIEW

To measure the performance and scalability, several search-centric tests were executed to show that operation p scaling out on seven nodes in OCI virtual compute shapes.

Test Flow Scenarios

Total number of users in LDAP directory – two million (2M)

Test duration – 30 minutes.

### OUD Search Scenarios:

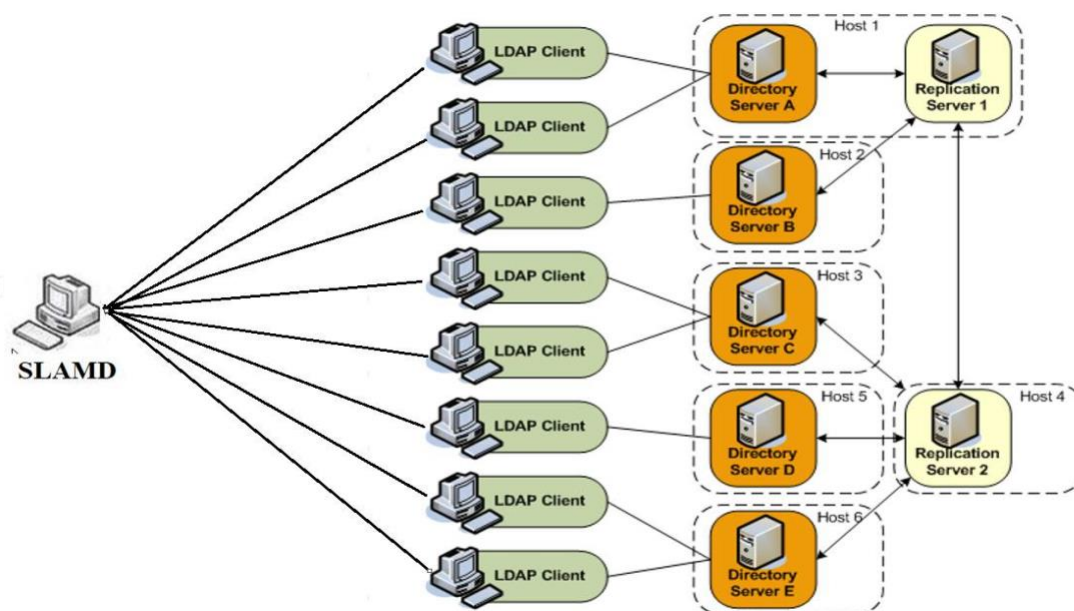
Test Case Name	Search Filters Pattern	Total Count
UID Exact Search	(uid=<id>)	10
UID Substring Search	(&(uid=<xx.*>)(<colmn>=<val>))	171
Complex Filter Substring Search with UID Exact Search	(&(uid=<id>)(telephoneNumber=*xxx*)(<title>=<*xx>))	10
Complex Filter Substring Search with UID Substring Search	(&(telephoneNumber=*xxxxx*)( (<title>=<*xx>)(uid=(x*>)))	378
Complex Filter Substring Search With Mail Substring Search	(&(telephoneNumber=*xxxxx*)( (<mail>=<xxxx.*>)(title=<*xx>)))	330
Complex Filter Substring with GivenName Substring Search-1	(&(telephoneNumber=<*xxxxx*>)( (<givenName>=<val>)(title=<*xx>)))	330

### Combined Load Scenarios:

Operations	load percentage
Bind with random user	20%
Search	59%
Compare operations	10%
Add operations	1%
Modify operations	10%

## Topology

Refer to the schematic picture below, describing the test configuration with all required OUD tiers, collocated in the same network segment to eliminate network latencies in the scalability runs.



## Hardware Configuration

Depending on the number of concurrent users, configurations can be categorized as small, medium, and large implementations. This section covers recommended hardware shapes, offered in OCI for each size to accommodate for the concurrent load, and projects for future growth. The customers can use the OCI shape specifications as the guidance for deploying comparable on-premises OUD configurations. Refer to the following link for more details on OCI compute shapes: <https://docs.oracle.com/en-us/iaas/Content/Compute/References/computeshapes.htm>

Table 1-1: OUD Node Hardware configuration (Small, Medium, Large are H/W shapes used for the sizing activity)

Configuration	Small	Medium	Large
Processor	AMD E3 - EPYC 7742 64-Core Processor		
Clock speed	2245.780GHz	2245.780GHz	2245.780GHz
OCI VM shape	VM.Standard.E3.Flex x	VM.Standard.E3.Flex	VM.Standard.E3.Flex
OCI shape RAM, Gb	16 - 64	32 - 64	64 +
OCI shape OCPUs	4	4 - 8	8 - 16

**Important considerations:** Server compute sizing specifications apply to a single node only. You can deploy another OUD compute node, using the same shape / specifications to scale-out and run as a cluster. When CPU utilization reaches 60-70% or higher on an Server node, you may have to plan for scaling out the cluster by adding another OUD node.

The recommended specifications are for a single database instance. Make sure you closely monitor the database load when you add more OUD Compute nodes in your environment.

## ATOMIC SEARCH BENCHMARKS

This chapter provides the results of the OUD load using atomic search operations for 2M LDAP users for small, medium, and large configurations. The benchmark was targeted to push the load till CPU reach ~90% and obtain the number of searches to get max operations per second and the number of concurrent LDAP connections. Note that for these tests 64Gb RAM nodes were used to run the LDAP directory size of two million in 'small', 'medium' and 'large' shapes.

Table 2-1: UID Exact Search

DEPLOYMENT TYPE	COMPUTE SHAPE	UID EXACT SEARCH		COMPLEX FILTER WITH EXACT SEARCH	
		LDAP connections	Max operations per second	LDAP connections	Max operations per second
Single node	Small	750	62,101	800	60,655
	Medium	1500	1,13,901	1800	1,05,754
	Large	3000	1,18,753	4000	1,16,233
7 node cluster	Small	5,300	4,65,864	6,000	4,23,312
	Medium	10,500	7,61,114	12,000	6,98,329
	Large	21,000	8,68,561	28,000	8,54,135


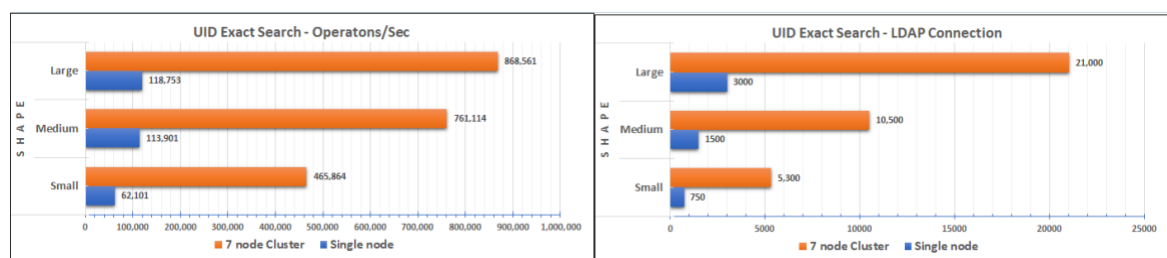
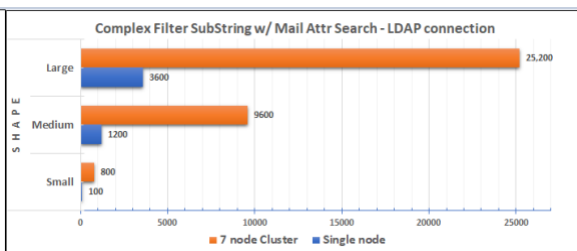
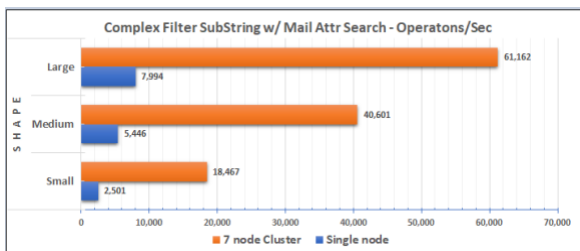




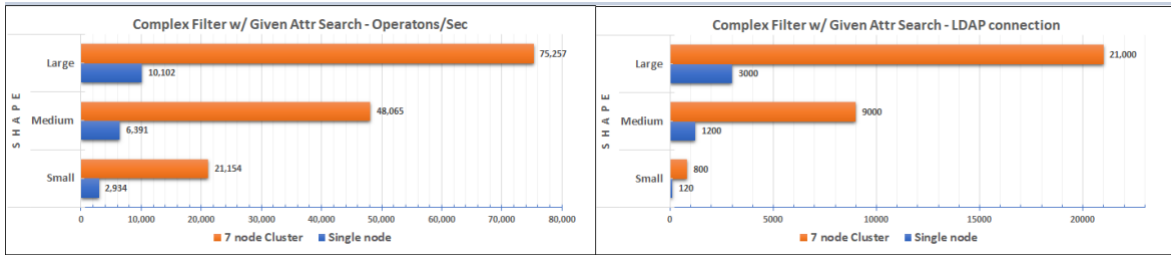


Table 2-2: Complex filter substring search

DEPLOYMENT TYPE	COMPUTE SHAPE	UID SUBSTRING SEARCH		COMPLEX FILTER W/ SUBSTRING SEARCH		COMPLEX FILTER SUBSTRING W/ MAIL ATTR SEARCH		COMPLEX FILTER W/ GIVEN ATTR SEARCH	
		LDAP connections	Max operations per sec	LDAP connections	Max operations per sec	LDAP connections	Max operations per sec	LDAP connections	Max operations per sec
single node	Small	100	1,661	100	2,740	100	2,501	120	2,934
	Medium	200	3,324	1250	6,315	1200	5,446	1200	6,391
	Large	1200	4,862	3000	9,976	3600	7,994	3000	10,102
7 node cluster	Small	400	12,616	800	19,897	800	18,467	800	21,154
	Medium	800	23,972	9,000	45,528	9,600	40,601	9,000	48,065
	Large	9,000	37,431	21,000	97,833	25,200	61,162	21,000	75,257



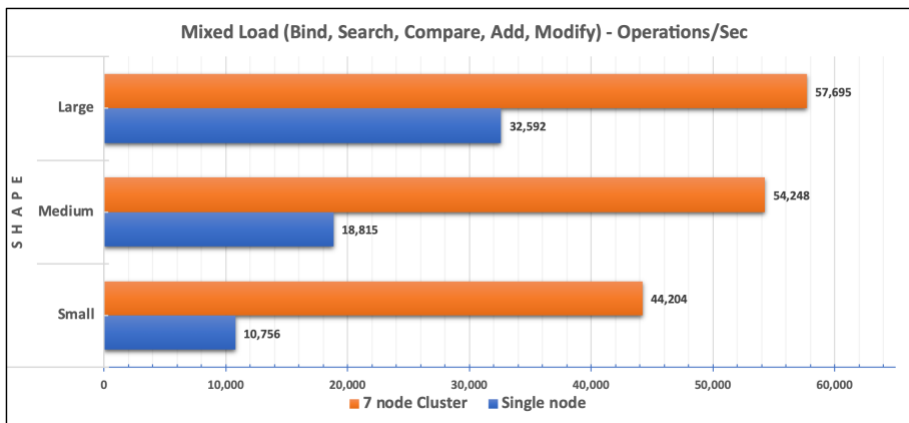


## COMBINED SEARCH BENCHMARKS

This chapter provides the results of OUD using combined search operations given in the “Test Flow Scenarios” section for 2M users with 1000 LDAP connections for small, medium, and large configurations.

Table 3-1: Mixed load result

Deployment type	Compute shape	Max operations/Sec
single node	Small	10,756
	Medium	18,815
	Large	32,592
7 node cluster	Small	44,204
	Medium	54,248
	Large	57,695



## PERFORMANCE TUNING RECOMMENDATIONS

The performance tuning recommendations are common for all configurations. Shape specific recommendations are referenced explicitly throughout the section.



## OID Tuning

**Table 4-1: Recommended configuration settings (default in OUD since April BP-220405)**

Tuning Parameters	Values
Multiple loggers disabled for new instances:	FALSE
File-Based Access Logger	FALSE
File-Based HTTP NSCA Access Logger	FALSE
File-Based HTTP Admin Logger	FALSE
File-Based Admin Logger	FALSE
Oracle Access Logger	FALSE
Oracle Admin Access Logger	FALSE
Oracle Error Logger	FALSE
Enable user entry cache	Enabled-FIFO
ds-cfg-tombstone-purge-interval	1d
ds-cfg-db-evictor-lru-only	FALSE
ds-cfg-db-evictor-nodes-per-scan	100
ds-cfg-db-checkpointer-bytes-interval	500 megabytes
ds-cfg-db-logging-level	WARNING
ds-cfg-db-num-cleaner-threads	10
ds-cfg-disk-low-threshold	200 megabytes
ds-cfg-disk-full-threshold	100 megabytes
ds-cfg-je-property: je.cleaner.lookAheadCacheSize=2097152	2097152(2mb)
ds-cfg-je-property: je.cleaner.bytesInterval=100000000	100000000(100mb)
ds-cfg-cache-level	Entry cache is at higher level than group cache.create-entry-cache --type fifo --set enabled:true --cache-name fifoEntryCache --set cache-level:1

Table 4-2: Additional recommended configuration settings by shape

Tuning Parameters	Small	Medium	Large
Default number of worker threads	Let the server decide	100	100
ds-cfg-entry-cache-preload	FALSE	FALSE	TRUE
ds-cfg-index-entry-limit	4000	10000	20000
ds-cfg-db-cache-percent	Let the server decide	60	60
ds-cfg-db-cleaner-min-utilization	50	50	75
ds-cfg-preload-time-limit	0ms	0ms	180s
group-membership-caching-threshold	Group membership will be cached only for those groups whose members are less than or equal to this specified threshold. Default value is 0 which means that no group members are cached. This parameter is applicable to all FIFO Group Caches for a given backend		

Note: **group-membership-caching-threshold** has been added to support better performance.

Updated the configuration file `/etc/security/limits.conf`:

```
opc    hard  nofile  65535
opc    soft  nofile  65535
root   hard  nofile  500000
root   soft  nofile  500000
#*     hard  nproc   500000
#*     soft  nproc   500000
```

#### OLD JVM Tuning

The table below provides the recommended heap size settings using sample LDAP directory sizes. Smaller directories running on nodes with less RAM can use the appropriate values from the recommended ranges below.

Table 4-3: Shape specific settings recommendations

Configuration	Small	Medium	Large
Sample LDAP directory size	5K	50K	2M
OCI shape RAM, Gb	8-16	16-32	64+

OID JVM Min Heap Size, Gb	4	8	16
OID JVM Max Heap Size, Gb	8	20	30+

In the above tests, the minimum and maximum heap sizes for OUD JVM were set to the higher values across all the shapes to accommodate for the benchmarked directory size of two million users.

Edit the instance-dir/OUD/config/java.properties file and set the following values. Run the dsjavaproperties command: bin/dsjavaproperties

```
start-ds.java-args="Xms30032m -Xmx30032m -d64 -server -Xmn4g -
XX:MaxTenuringThreshold=1 -XX:+UseConcMarkSweepGC -
XX:CMSInitiatingOccupancyFraction=70 -XX:+UseCMSInitiatingOccupancyOnly -
XX:+CMSScavengeBeforeRemark -XX:+UseCompressedOops"
```

**Note:** -XX:+USECOMPRESSEDPOOPS was set to enable compressed pointers (object references represented as 32-bit offsets instead of 64-bit pointers) for optimized 64-bit performance with Java heap sizes less than 32gb.

## CONCLUSION

This document provides actionable recommendations for sizing Oracle Unified Directory 12.2.1.4.0 to help ensure best performance and scalability.


## GLOSSARY

- IDM – Oracle Identity Management
- IAM – Oracle Identity and Access Management
- OUD – Oracle Unified Directory
- RAC – Oracle Real Application Cluster
- RT – Response Time
- OPS – operations per seconds

---

Connect with us

Call +1.800.ORACLE1 or visit [oracle.com](https://www.oracle.com). Outside North America, find your local office at: [oracle.com/contact](https://www.oracle.com/contact).

 [blogs.oracle.com](https://blogs.oracle.com)

 [facebook.com/oracle](https://www.facebook.com/oracle)

 [twitter.com/oracle](https://twitter.com/oracle)

---

---

Copyright © 2023, 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.

This device has not been authorized as required by the rules of the Federal Communications Commission. This device is not, and may not be, offered for sale or lease, or sold or leased, until authorization is obtained.

---

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. 0120

Disclaimer: If you are unsure whether your data sheet needs a disclaimer, read the revenue recognition policy. If you have further questions about your content and the disclaimer requirements, e-mail [REVREC\\_US@oracle.com](mailto:REVREC_US@oracle.com).