An Oracle White Paper
March 2013

Oracle Internet Directory 11g Benchmark With 50 Million Users On Oracle SPARC T5-2 Server

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.

Introduction

In the following whitepaper we are focusing on the specific performance improvements related to the entry cache of Oracle Internet Directory (OID) 11.1.1.7. Oracle Internet Directory was picked as one of the applications to showcase the performance of the newly released SPARC T5 based line of servers, and how the performance numbers stack up to other hardware architectures.

The new entry cache now resides in shared memory, so multiple Oracle Internet Directory server instances on the same host can share a cache. Previous versions of OID only supported an entry cache for one and only one OID instance, even though for high availability and enhanced performance multiple OID instances have been deployed. This limitation is now lifted, and we can see improved OID server performance of up to 300% compared to the previous implementation.

With impressive 944,624 search operations / sec on a single SPARC T5-2 server, with 50M user entries, Oracle Internet Directory demonstrates linear server performance from 2 to 32 cores on a single SPARC T5-2 server. OID's unique software architecture provides unmatched scalability per server node, and across multiple servers[1] to deliver performance on demand.

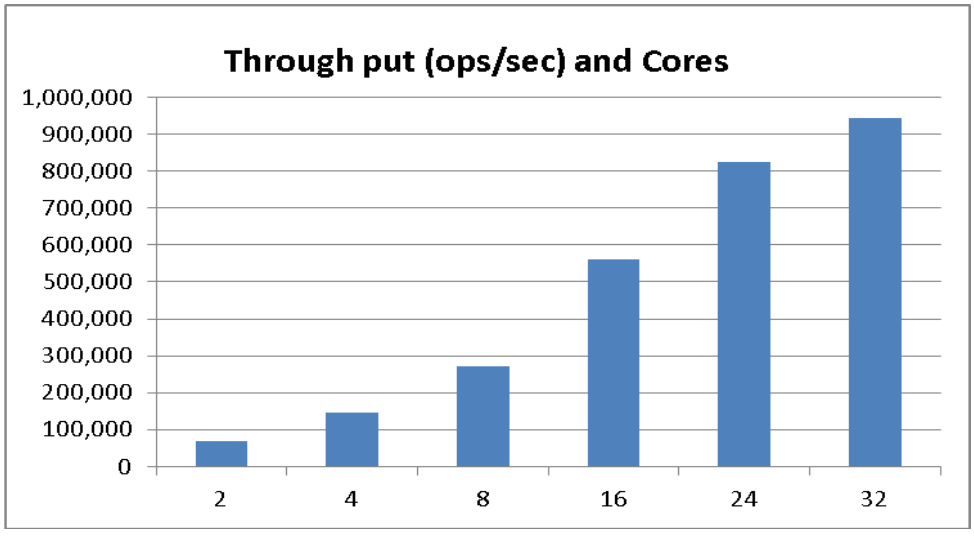

Figure 1: Oracle Internet Directory Scalability

[1] **Oracle Internet Directory 11g Benchmark With 10 Million Users On Oracle Exalogic X2-2**

The highlights of the benchmark are summarized below –

- 50M Entries loaded into OID under 4 hours.
- 944,624 LDAP searches /sec, average latency of 1.05 msec with 1000 clients.
- 305,307 authentications /sec, average latency of 3.27 msec with 1000 clients.
- 594,426 comparisons /sec of userpassword attribute, average latency of 1.68 msec with 1000 clients.
- 16,735 LDAP modifies /sec with an average latency of 2.98 msec with 50 clients.
- 1000 LDAP additions /sec with an average latency of 15.95 msec with 16 clients.

All operations were performed on random entries and with ACLs defined.

Oracle Internet Directory Details

Oracle Internet Directory (OID) implements a unique architecture which enables the directory to fully utilize the underlying server hardware, scale on any given hardware, and at the same time provide high availability.

This architecture

- Multi-threaded using DB connection pooling
- Multi-process to utilize all available CPU's
- Multi-instance directory server using multiple HW nodes
- Scalability with the number of CPUs in SMP HW architectures
- Scalability with the number of nodes in HW cluster architectures

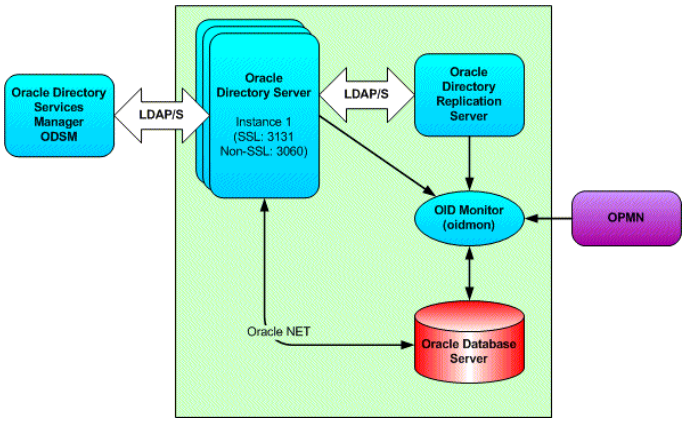

Figure 2: Oracle Internet Directory typical node architecture

In a typical deployment one or more OID servers, together with the OID replication server, are usually co-located with the DB instance on the same physical host. For more details regarding OID monitor and Oracle Process Management and Notification (OPMN) see the OID Admin guide.

For further details regarding Oracle High Availability please refer to the Fusion Middleware High Availability Guide and the Identity Management Deployment Guide. Both outline the recommended deployment architectures and serve as a blueprint for an Enterprise deployment.

Oracle SPARC T5-2 Overview

The new SPARC T5-8, T5-4, T5-2 and T5-1B servers are based on the SPARC T5 processor, the world's fastest microprocessor. The SPARC M5-32 is based on the SPARC M5 processor, which shares the same core as the SPARC T5 processor. These systems enable great acceleration of critical business workloads, extraordinary scalability and capacity, and leverage Oracle Solaris 11, the best UNIX operating system available today.

The new SPARC servers have shattered 17 performance records. The SPARC M5-32 server has shown it can be up to 10X faster than its previous generations. All Oracle customers – whether they are growing server infrastructure for enterprise applications, or consolidating existing server through virtualization technologies – can benefit from the exceptional speed and reliability of the new SPARC T5 and M5 servers.

With the release of these new servers, Oracle offers the most powerful and efficient SPARC-based server family ever. The SPARC M5-32 server is Oracle's new offering at the high-end, while the SPARC T5 servers represent Oracle's new mid-range line-up. The SPARC T4 servers complement the SPARC server line at the entry-level. Based on SPARC M5, SPARC T5, and SPARC T4 processors—which all share the same processor core—the SPARC-based server family provides seamless scalability from 1 up to 32 processors.

All of the servers in the SPARC-based family are designed with mission-critical applications in mind, and run the Oracle Solaris operating system—the best UNIX system for Oracle deployments. They share the same virtualization capabilities through Oracle VM Server for SPARC and leverage the same systems management framework through Oracle Enterprise Manager Ops Center. This leads to unprecedented simplicity in the deployment of all enterprise workloads, enabling reduction of business risk, delivering savings in management costs, and unlocking flexibility to grow your business to any scale, while maximizing reliability and uptime.

Benchmark Deployment Architecture

From a high level perspective the setup is outlined in Figure 3, using a two tier deployment.

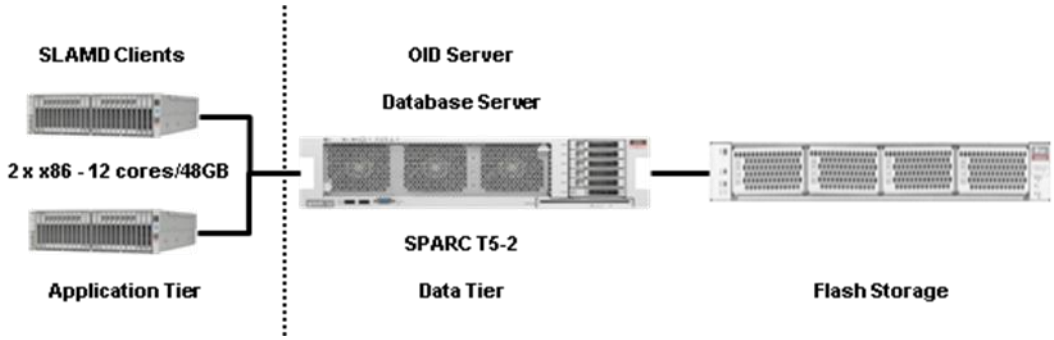

Figure 3. Oracle Internet Directory Deployment Architecture Overview

To run the benchmark we used the SLAMD distributed load generation engine. SLAMD clients were distributed across two X4170 nodes (Application Tier) and connected to the OID server instances on a single SPARC T5-2 server. The backend OID database was run on the same SPARC T5-2 server (Data Tier).

All systems were inter-connected via 10GB network.

Hardware and OS Configuration

OID/DB Hardware details

OID and database co-existed in the same hardware.

CPU	Oracle SPARC T5-2 (Total 32 cores), 3.6GHZ, 8M L3 cache
RAM	512GB
IO	2 X Internal disks 2 X 10 Gigabit Ethernet card 1 X 1 GBE card
Storage	Sun Storage F5100 Flash Array
OS	Oracle Solaris 11.1

Client Node details

SLAMD tool was used to generate the LDIF file and run the benchmark test. We used two nodes for slamd clients. Each client node consists of

CPU	2 Six-Core Intel CPU
RAM	48GB
IO	1 X Internal disks 2 X 10 Gigabit Ethernet card 1 X 1 GBE card
OS	OEL5

Database Tuning

Out of the box configuration was retained with the following modifications:

Parameter	Value	Description
SGA_MAX_SIZE	96GB	The hard limit up to which sga_target_size can dynamically adjust sizes for the System Global Area (SGA).
SGA_TARGET	96GB	The target size for the System Global Area.
PGA_AGGREGATE_TARGET	16GB	Target size for the Program Global Area.
AUDIT_TRAIL	none	Configures database auditing.
Processes	5000	Number of database processes.
FILESYSTEMIO_OPTIONS	setall	
_*Disk*_sector_size_override	TRUE	needed if redo is used on flash disk
_simulate_disk_sectorsize	4096	needed if redo is used on flash disk
DB_FILE_MULTIBLOCK_READ_COUNT	128	

Database Redo Logs

We created the REDO logs with 4K-block size; this is required for SSD drives only.

OID Tuning

Out of the box OID configuration was retained for this benchmark. This means that all default access controls and password policies were effective and in play. However, to best utilize the hardware resources at hand minimal tuning was performed by changing the following OID configuration parameters.

Parameter	Value	Description
ORCLSERVERPROCS	32	Number of server processes
ORCLMAXCC	8	Number of worker threads per server process. This was the optimal configuration for this hardware.
ORCLECACHEMAXSIZE	260GB	Maximum size of the entry cache. This was based on the available memory in the hardware.
ORCLECACHEMAXENTRIES	50 Million	Maximum number of entries that can be cached. This was calculated based on entry size and cache size.
ORCLEXTCONFFLAG	0	Disabled auditing logs.
ORCLECACHEENABLED	2	Enable entry cache & Result set cache
ORCLSKIPREFINSQL	1	Skip referral processing in SQL, since there are no referral entries in this test.
ORCLMATCHDNENABLED	0	Disabled match DN, this does not affect successful search.
OID_BIND_CPU	24	This enables OID to create processor set and bind to it. By setting OID_BIND_CPU to 24 OID creates 24 processor set (one per CPU core) and binds 32 server process to 24 processor set. We need to have 8 CPU cores for database processing.

Data Characteristics

Oracle Internet Directory (LDAP Data)

50 million LDAP user entries were generated using the SLAMD makeLDIF template shown in Appendix A. These entries had the following properties:

- Total 36 attributes
- 12 generated operational attributes.
- Standard InetOrgPerson entry.

Finally, it is important to note that no data partitioning was necessary to accommodate the workload scenario.

Oracle Internet Directory Database & Bulkload

PRE-BULKLOAD TUNING

Below steps are required before bulkload

- Set PGA_AGGREGATE_TARGET=64G in database.
- Create temporary tablespace of 100G.
- Edit OID ORACLE_HOME/ldap/admin/oidtblkl.sql in procedure crtidx(), change PARALLEL to NOPARALLEL to avoid spawning too many Oracle process during index creation. Bulkload is anyway performing parallel index of different tables.

Resize OID tablespace for bulkload/add/modify tests

It is recommended that pro-active resizing of OID tablespace before running bulkload/add/modify tests.

Below is the sql command to increase OID tablespace.

```
Sqlplus ods@oiddb  << eof
alter tablespace olts_attrstore resize 180g;
alter tablespace olts_ct_store resize 150g;
eof
```

Bulkload loading the data

```
bulkload connect=oiddb check=true generate=true threads=16
file=<PathToLdif>
```

Note: threads=16 is very important otherwise bulkload will spawn 256 threads which is not optimal. This is an offline activity.

Make sure no errors reported then invoke load

```
bulkload connect=oiddb load=true threads=16 file=<pathToLdif>
```

Bulk Load Timings

- Bulk Load “generate” phase = 1 hour 20 min 41 sec, we used just generate=true option.
- Bulk Load “load”, “index” and “Statistics” creation = 2 hours 48 minutes

The database representation of this data had the following characteristics:

TableSpace Name	Size in GB
OLTS_ATTRSTORE	134
OLTS_CT_STORE	116
UNDOTBS1	17
OLTS_DEFAULT	1.6
REDO_LOGS	16 X 3

Workload Scenario

The workload scenario tested was a 50 million user deployment with exhaustive load. This scenario examines the case of a large deployment where all the users in the deployment are active at a given time. Therefore in our tests benchmarking this scenario we targeted all 50 million users deployed.

Test Scenarios

Each of the following read tests were executed with the above workload scenario.

Random LDAP Search Operations Test

This test scenario involved concurrent clients binding once to OID and then performing repeated LDAP Search operations. The salient characteristics of this test scenario is as follows

- SLAMD 'SearchRate' job was used.
- BaseDN of the Search is root of the DIT, the scope is SUBTREE, the search filter is of the form "UID=<a unique value>", DN is the required attribute to be returned.
- Each LDAP search operation matches a single entry
- The total number of concurrent clients was 1000 and were distributed amongst client nodes
- Each client binds to OID once and performs repeated LDAP Search operations, each Search operation resulting in the lookup of a unique entry in such a way that no client looks up the same entry twice and no two client's lookup the same entry and all entries are searched **randomly**.
- Test job was run for 60 minutes.

Random LDAP Authentication Operations Test

This test scenario involved concurrent clients repeatedly executing the sequence of performing an LDAP Search operation to look up an user and performing a simple bind as that user to verify its credential. The salient characteristics of this test scenario is as follows –

- SLAMD 'AuthRate' job was used.
- BaseDN of the Search is root of the user container, the scope is BASE, the search filter is of the form "UID=<a unique value>", DN is the required attribute to be returned.
- Each LDAP search operation matches a single entry

- All entries had the same userpassword value
- The total number of concurrent clients was 1000 and were distributed amongst client nodes
- Each client binds to OID once and performs repeated LDAP Search followed by bind operations where each Search operation results in the lookup of a unique entry in such a way that no client looks up the same entry twice and no two clients lookup the same entry and all entries are searched **randomly**.
- Test job was run for 60 minutes.

Random LDAP Compare Operations Test

This test scenario involved concurrent clients binding once to OID and then performing repeated LDAP Compare operations. The salient characteristics of this test scenario is as follows –

- SLAMD 'CompareRate' job was used.
- BaseDN of the Search is root of the DIT.
- Each LDAP Compare operation, compares userpassword attribute
- The total number of concurrent clients was 1000 and were distributed amongst client nodes
- Each client binds to OID once and performs repeated LDAP Compare operations, each Compare operation resulting in the lookup of a unique entry in such a way that no client compare the same entry twice and and all entries are compared **randomly**.
- Test job was run for 60 minutes.

Random LDAP Modify Operations Test

This test scenario involved concurrent clients binding once to OID and then performing repeated LDAP Modify operations. The salient characteristics of this test scenario is as follows –

- SLAMD 'ModRate' job was used.
- Each client updates a unique entry each time and a total of 50 Million entries are updated
- The total number of concurrent clients was 50 and were distributed amongst client nodes

- Each client binds to OID once and performs repeated LDAP Modify operations, and all entries are modified **randomly**.
- Test job was run for 30 minutes.

LDAP Mixed Load Test

This test scenario involved concurrent clients binding once to OID and then performing repeated LDAP Mix operations. The salient characteristics of this test scenario is as follows

- The ratio involved 65% LDAP Search clients, 30% LDAP Bind and 5% LDAP Modify clients.
- A total of 1000 concurrent LDAP clients were used and were distributed on 2 client nodes.
- Each client binds to OID once and performs repeated LDAP operations, and all entries are modified/lookedup **randomly.**
- Test job was run for 30 minutes.

LDAP Add Load Test

This test scenario involved concurrent clients binding once to OID and then performing repeated LDAP Mix operations. The salient characteristics of this test scenario is as follows

- A total of 500,000 entries were added.
- A total of 16 concurrent LDAP clients were used.
- Slamd add's 'inetorgperson' objectclass entry with 21 attributes (includes operational attributes).

Benchmark Results

A summary of the results is tabulated in the below table.

Test Scenario	Number of Clients	Through put (ops/sec)	Latency (msec)
SearchRate	1000	944624	1.05
CompareRate	1000	594426	1.68
AuthRate	1000	305307	3.27
Modify	50	16735	2.98
Mix **BIND** **MODIFY** **SEARCH**	1000 300 50 650	291360 87159 14528 188832	3.43 1.08 12.0 3.86
Add	16	1000	15.95

CPU SCALABILITY TEST (SPARC T4-4 AND T5-2)

CPU Cores	OID server process (orclserverprocs)	Slamd clients	Through put (ops/sec) T5-2	Latency (msec) T5-2	Through put (ops/sec) T4-4	Latency (msec) T4-4
2	2	100	68399	1.46	59895	1.6
4	4	100	145879	0.68	112258	0.8
8	8	500	270601	1.84	195263	2.5
16	16	500	560709	0.88	378971	1.3
24	24	1000	823741	1.21	546695	1.2
32	32	1000	944624	1.05	682133	1.4

SPARC T5-2 VS INTEL

Sparc T5-2, 3.6GHZ, 12 cores (used 12 cores from 32 cores)

Intel X5675, 3.06GHZ, 12 cores (2 X 6 cores)

Test	Slamd Clients	T5-2 (Ops/sec)	T5-2 latency (ms)	Intel (Ops/sec)	Intel latency(ms)
Messaging	500	417344	1.19	299724	1.66
Compare	500	274185	1.82	202433	2.46
Authentication	500	149623	3.3	119198	4.19

LDAP Search Operations (32 cores)

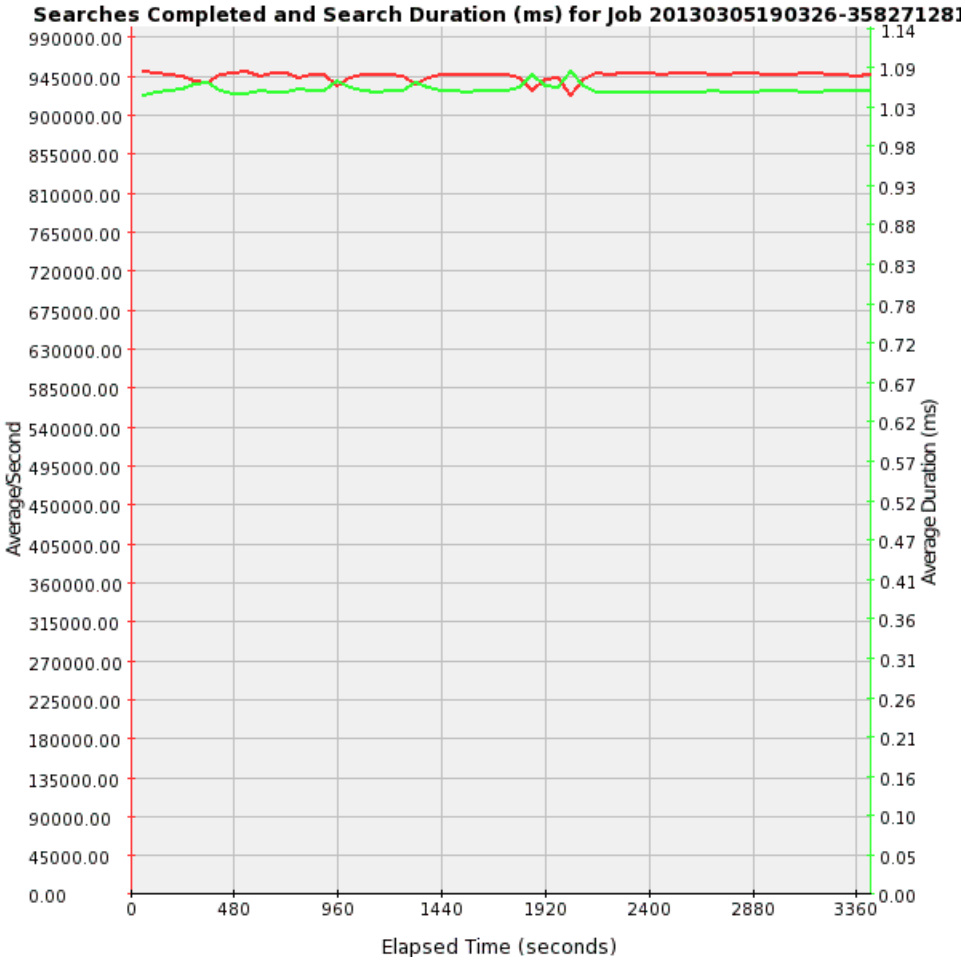

LDAP Compare Operations (32 cores)

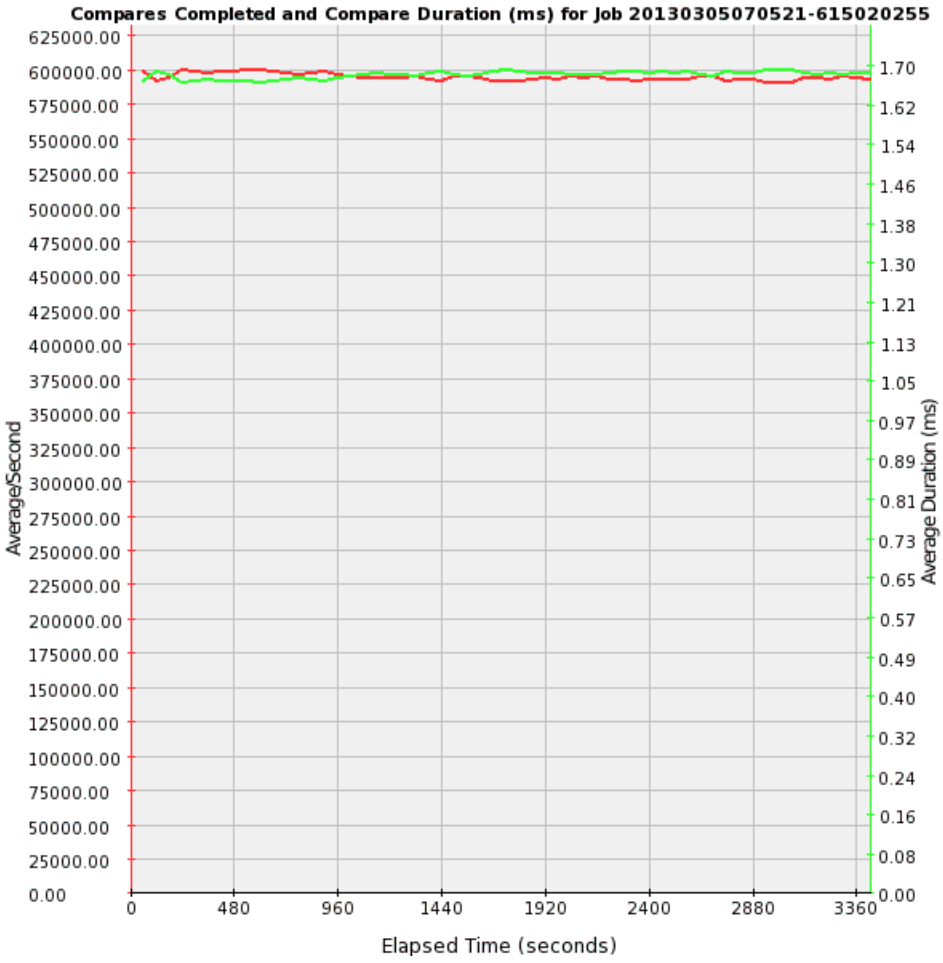

LDAP Authentication Operations (32 cores)

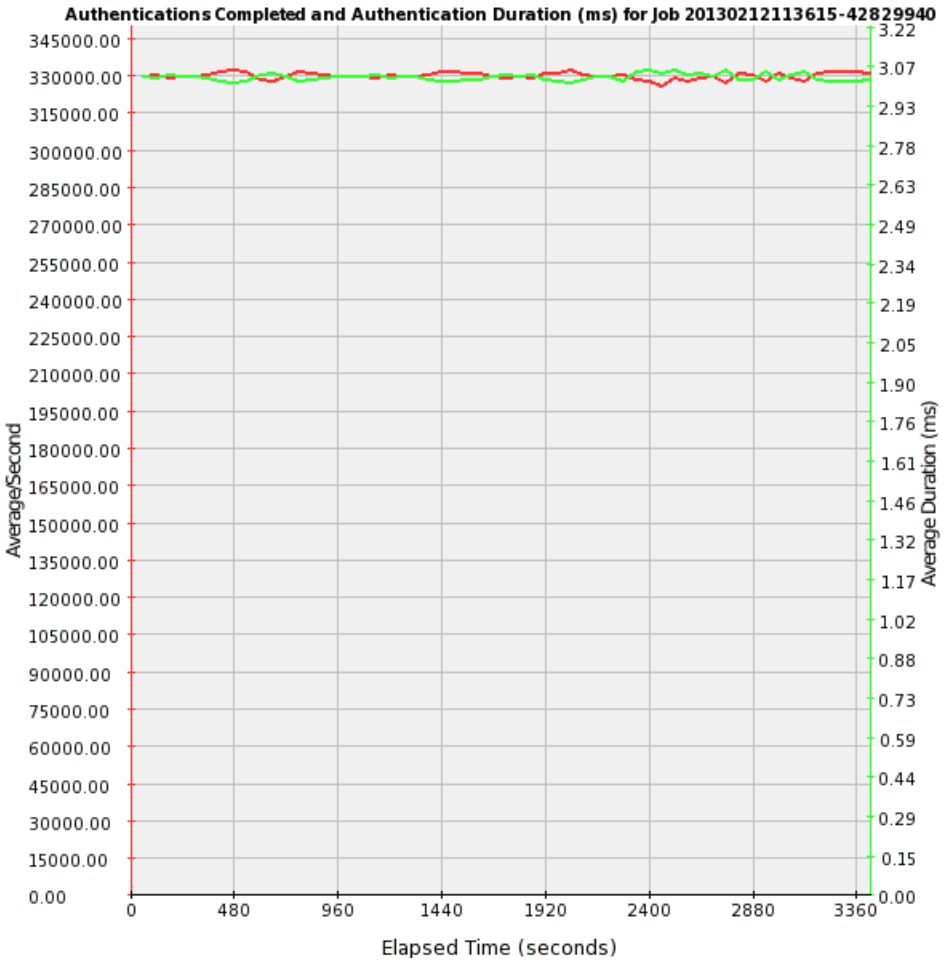

LDAP Modify Operations

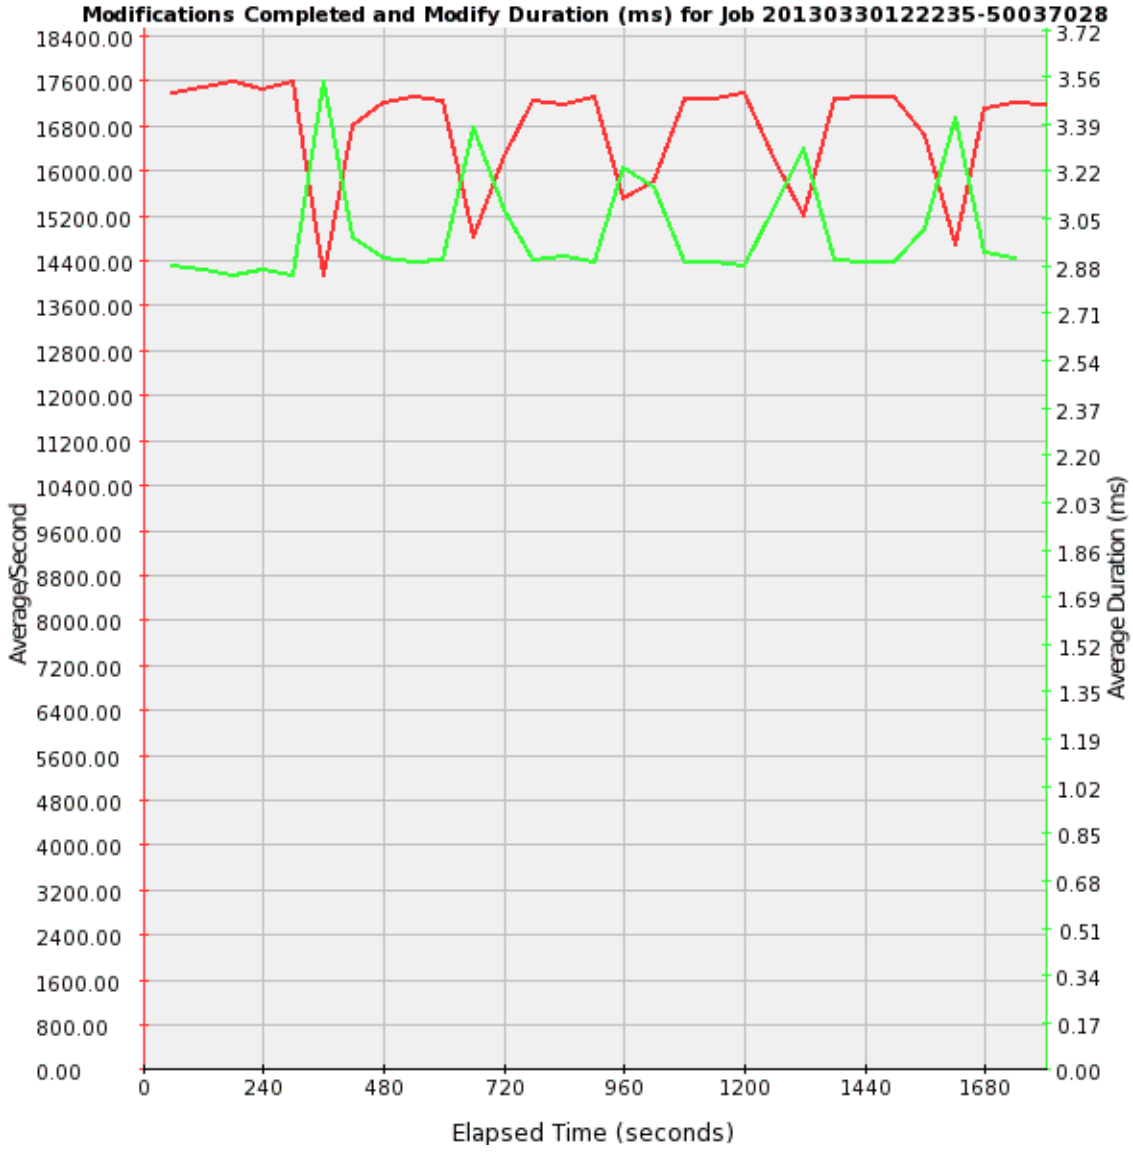

LDAP MIX Operations

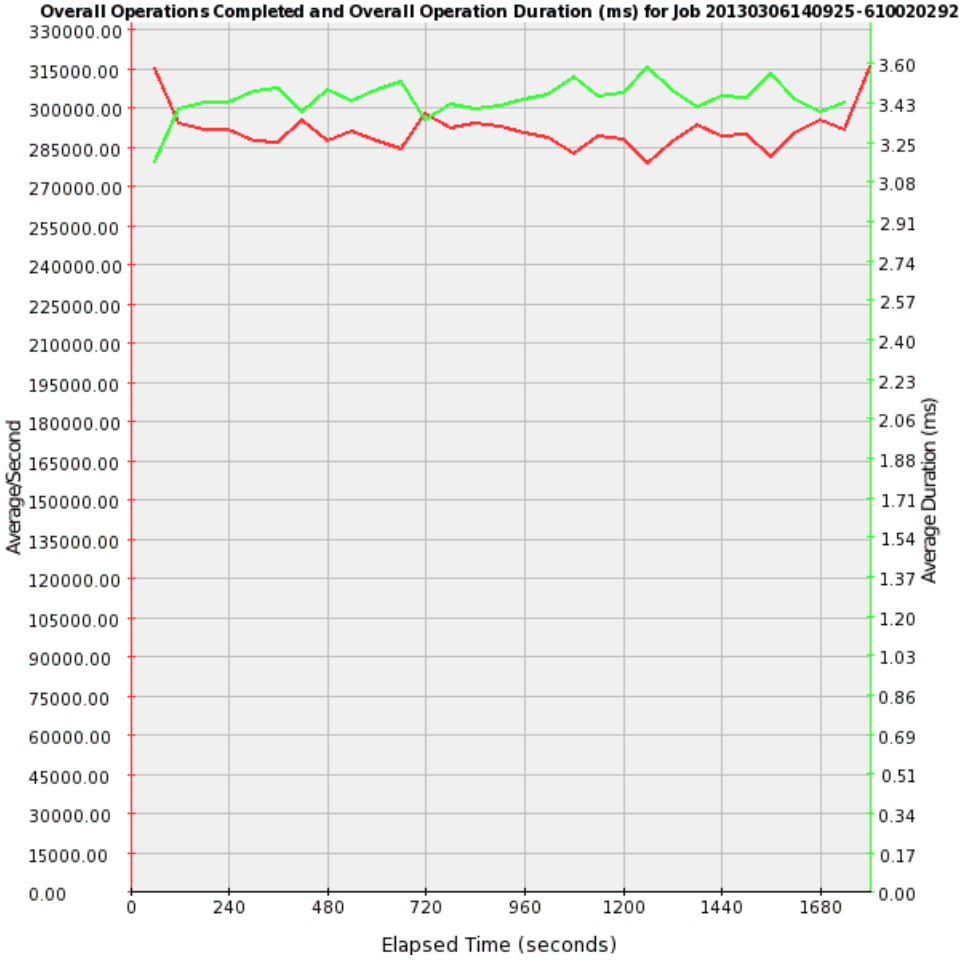

LDAP ADD Operation

LDAP Search Operations (24 cores)

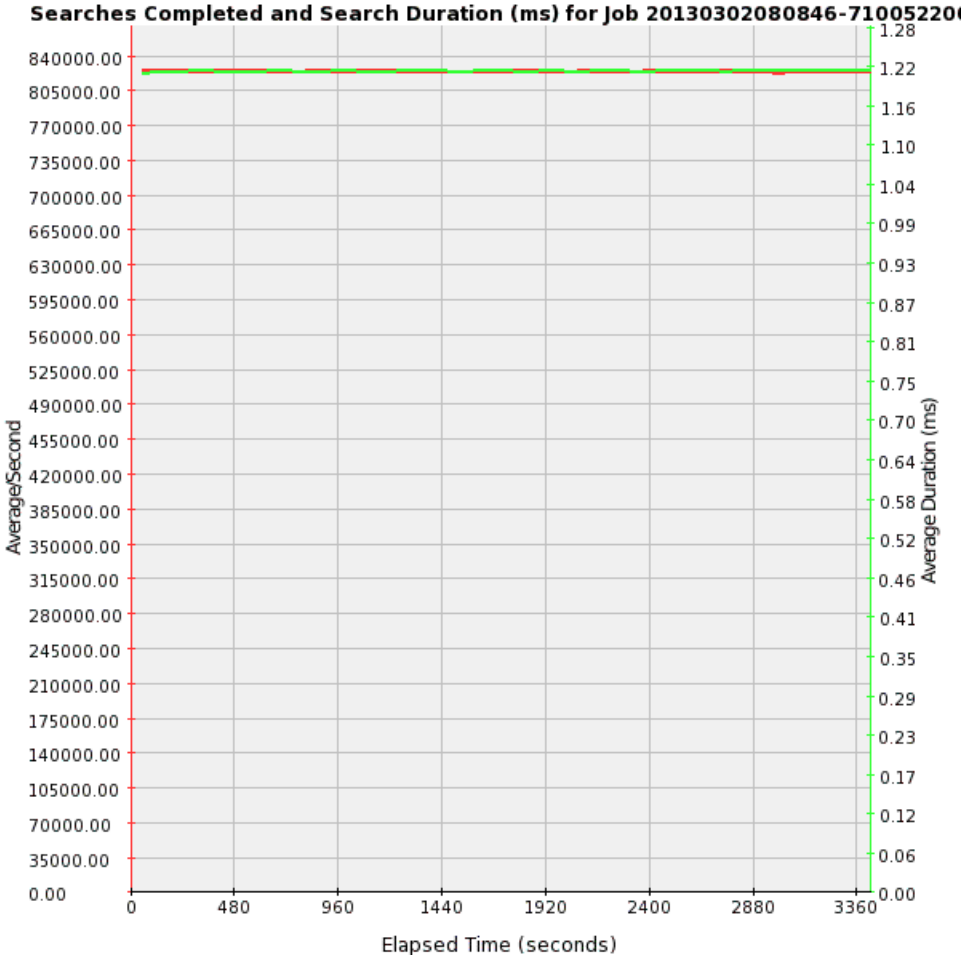

LDAP Search Operations (16 cores)

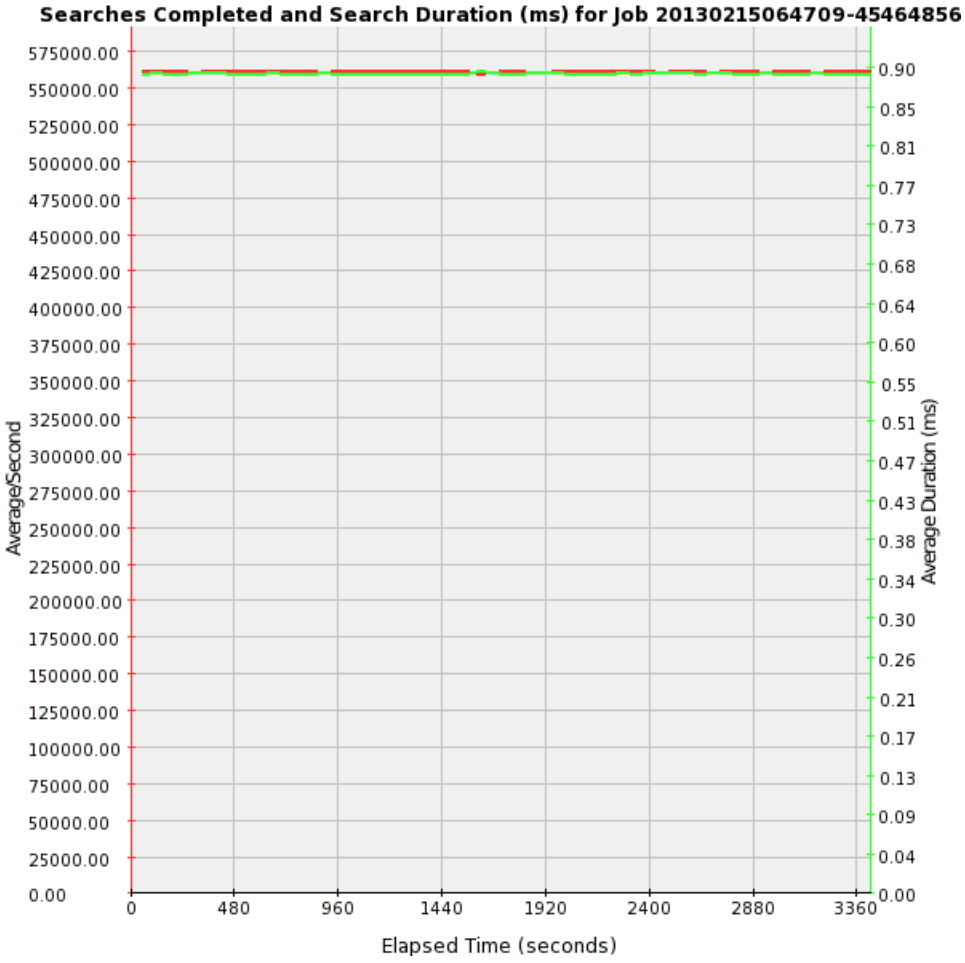

LDAP Search Operations (8 cores)

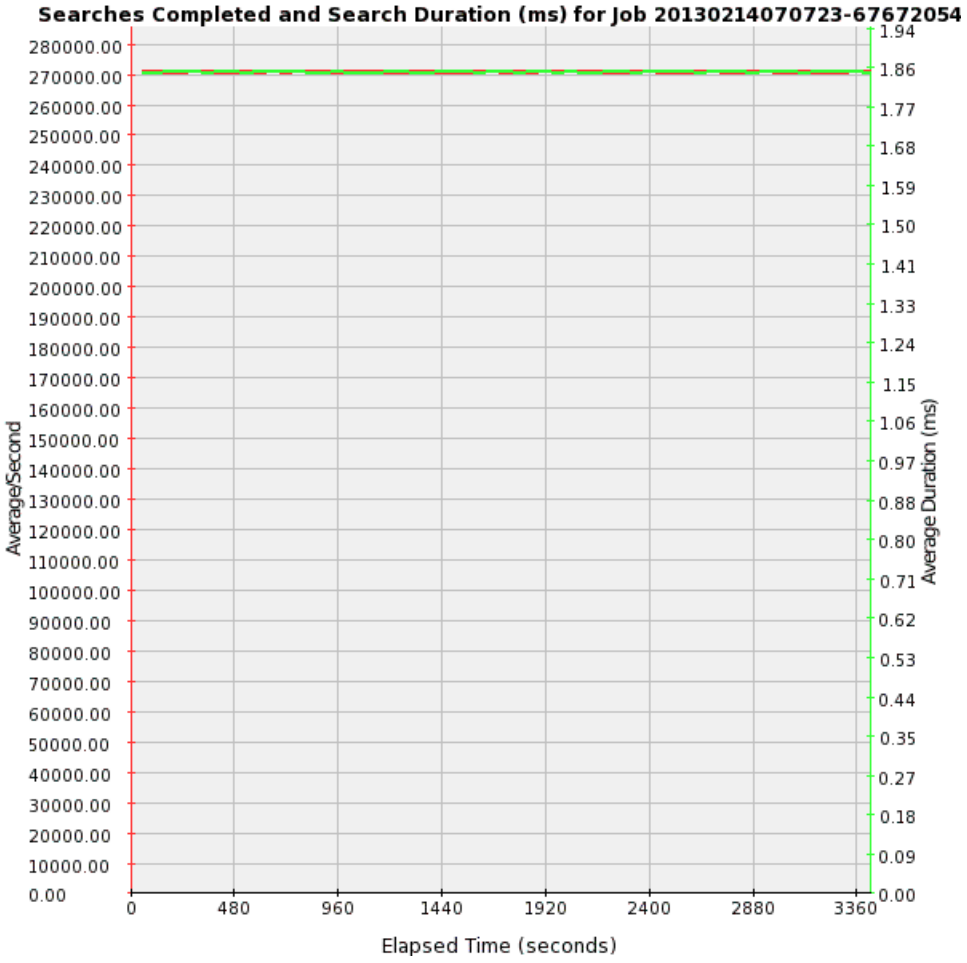

LDAP Search Operations (4 cores)

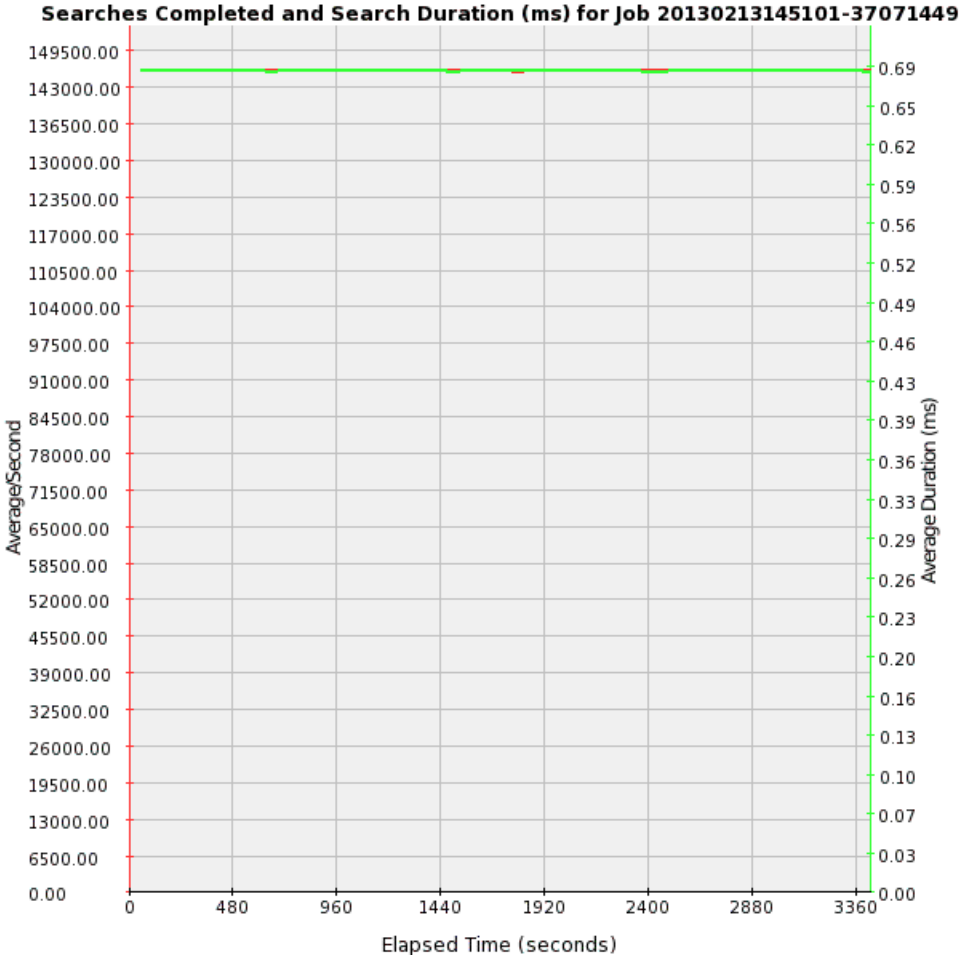

LDAP Search Operations (2 cores)

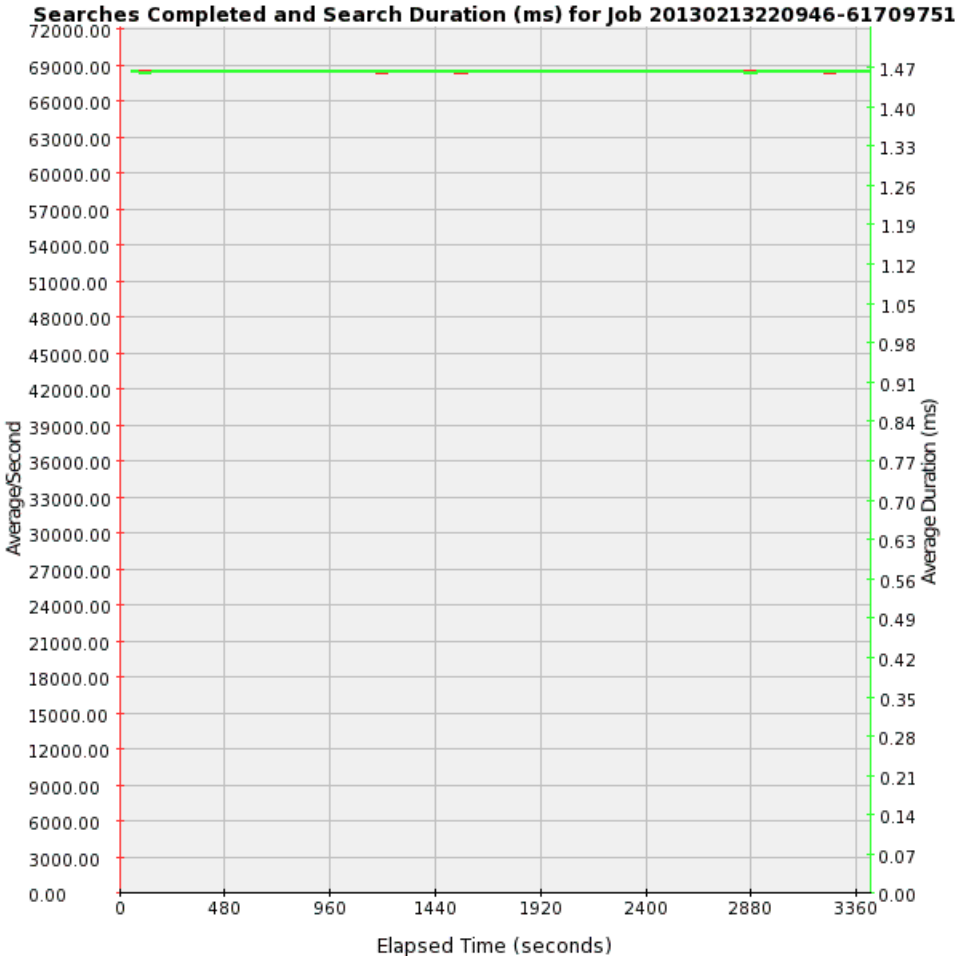

APPENDIX A

```
Slamd template file for generating the LDIF file

define suffix=ou=people,dc=us,dc=oracle,dc=com

define maildomain=oracle.com

define numusers=50000000

branch: ou=People,[suffix]

subordinateTemplate: person:[numusers]

-----------------------------------------------------------------------
-------

template: person

rdnAttr: uid

objectClass: top

objectClass: person

objectClass: organizationalPerson

objectClass: inetOrgPerson

givenName: <first>

sn: <last>

cn: {givenName} {sn}

initials: {givenName:1}{sn:1}

uid: user.<sequential:1>

mail: {uid}@[maildomain]

userPassword: password

telephoneNumber: <random:telephone>

homePhone: <random:telephone>

pager: <random:telephone>

mobile: <random:telephone>

employeeNumber: <sequential:1>

street: <random:numeric:5> <file:streets> Street

l: <file:cities>

st: <file:states>

postalCode: <random:numeric:5>

postalAddress: {cn}${street}${l}, {st}  {postalCode}

description: This is the description for {cn}.
```

APPENDIX B

Init.Ora parameters

```
orcl.__db_cache_size=90194313216
orcl.__java_pool_size=268435456
orcl.__large_pool_size=1879048192
orcl.__oracle_base='/export/home/oid/app/ram'#ORACLE_BASE set from
environment
orcl.__pga_aggregate_target=17179869184
orcl.__sga_target=103079215104
orcl.__shared_io_pool_size=0
orcl.__shared_pool_size=9663676416
orcl.__streams_pool_size=268435456
*.audit_file_dest='/export/home/oid/app/ram/admin/orcl/adump'
*.audit_trail='none'
*.compatible='11.2.0.0.0'
*.control_files='/OIDVOL1/oid/oradata/orcl/control01.ctl','/export/ho
me/oid/app/ram/fast_recovery_area/orcl/control02.ctl'
*.db_block_size=8192
*.db_domain=''
*.db_name='orcl'
*.db_recovery_file_dest='/export/home/oid/app/ram/fast_recovery_area'
*.db_recovery_file_dest_size=4322230272
*.diagnostic_dest='/export/home/oid/app/ram'
*.dispatchers='(PROTOCOL=TCP) (SERVICE=orclXDB)'
*.filesystemio_options='setall'
*.open_cursors=300
*.pga_aggregate_target=17179869184
*.processes=5000
*.remote_login_passwordfile='EXCLUSIVE'
*.sga_max_size=103079215104
*.sga_target=103079215104
*.undo_tablespace='UNDOTBS1'
_disk_sector_size_override=TRUE
_simulate_disk_sectorsize=4096
_fg_sync_sleep_usecs=1
_use_adaptive_log_file_sync=TRUE
```

APPENDIX B

Running with fixed priority for OIDLDAPD process

ps -ef | grep oidldapd | grep -v grep | awk '{print $2}' | xargs priocntl -s -c FX -m 50 -p 50 -i pid

For the update tests, set the priority for log writer and Oracle foreground process as shown below

priocntl -s -c FX -m 60 -p 60 -i pid <logWriterPID>

All other foreground process should run with priority 55

ps -ef | grep LOCAL=NO | grep -v grep | awk '{print $2}' | xargs priocntl -s -c FX -m 55 -p 55 -i pid

APPENDIX C

CPU usage can be controlled in OID by setting orclserverprocs parameter in the configuration entry "cn=oid1,cn=osdldapd,cn=subconfigsubenty". Environment variable OID_BIND_CPU should be set to 24. OID server automatically creates processor sets based on the value "orclserverprocs" and bind to it. Value of orclserverprocs should not be more than available CPU cores on the system.

Appendix D: Additional Information

1. Oracle Internet Directory 11g And Oracle Exadata In The Facebook Age
2. Oracle Internet Directory 11g benchmark on Exalogic X2-2
3. 2 Billion User Benchmark (Oracle Internet Directory 10.1.4.0.1)
4. Oracle® Fusion Middleware Enterprise Deployment Guide for Oracle Identity Management 11g Release 1 (11.1.1)
5. Oracle® Fusion Middleware High Availability Guide 11g Release 1 (11.1.1)
6. Oracle Internet Directory on the Oracle Technology Network
7. SLAMD Load Generation Engine

Oracle Internet Directory, 50 Million User Directory benchmark on SPARC T5-2
March 2013

Author: Ramaprakash Sathyanarayan
Contributors: Uday Shetty, Giri Mandalika, Olaf Stullich, Ravi Thammaiah

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com

Oracle is committed to developing practices and products that help protect the environment

Hardware and Software, Engineered to Work Together