50 Million User Directory Benchmark

Oracle Internet Directory 11.1.1.7.11 on Oracle's SPARC T7 Server

An Oracle White Paper
November 2015

EXECUTIVE OVERVIEW

Oracle Internet Directory (OID) is Oracle's LDAPv3 Directory Server that is highly scalable, available and manageable. OID has multi-threaded, multi-process, multi-instance process architecture with Oracle Database as the directory store. This unique physical architecture enables OID to take advantage of Oracle's SPARC T7 server running Oracle Solaris to achieve excellent performance and linearly scale to large number of cores.

INTRODUCTION

ORACLE INTERNET DIRECTORY

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 compared to the previous implementation.

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

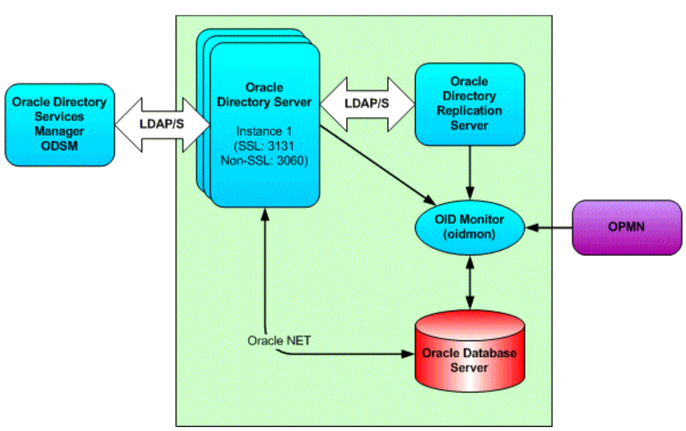

Figure 2: Oracle Internet Directory typical node architecture

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.

ORACLE SPARC M7/T7

The new SPARC T7 and SPARC M7 servers are based on the SPARC M7 processor, the world's fastest microprocessor. 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 SPARC M7 processor has support for Secure Silicon Memory and Data Acceleration at chip level providing Enterprise Security and In-Memory performance. 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 T7 and M7 servers.The SPARC M7 servers are Oracle's new offering at the high-end, while the SPARC T7 servers represent Oracle's new mid-range and low end line-up. The SPARC-based server family provides seamless scalability from 1 up to 16 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, Oracle Solaris Zones 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 Highlights

50M Entries were loaded in bulk into OID in under 2 hours.

Up to 1.17M LDAP searches/sec, with an average latency of 0.84 msec with 1000 clients was achieved.

Up to 414,394 authentication / sec with an average latency of 2.41 msec with 1000 clients was achieved.

Up to 650,047 compare of userpassword attribute with average latency of 1.53 msec with 1000 clients was achieved.

Up to 22,408 LDAP modify/sec with an average latency of 2.23 msec with 100 clients was achieved.

Up to 1436 LDAP add/sec with an average latency of 11.094 msec with 100 clients was achieved.

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

With impressive 1,177,947 search operations/sec on a single SPARC T7 chip, with 50M user entries, Oracle Internet Directory demonstrates linear server performance from 2 to 31 cores on a single SPARC T7 socket. OID's unique software architecture provides unmatched scalability per server node, and across multiple servers1 to deliver performance on demand.

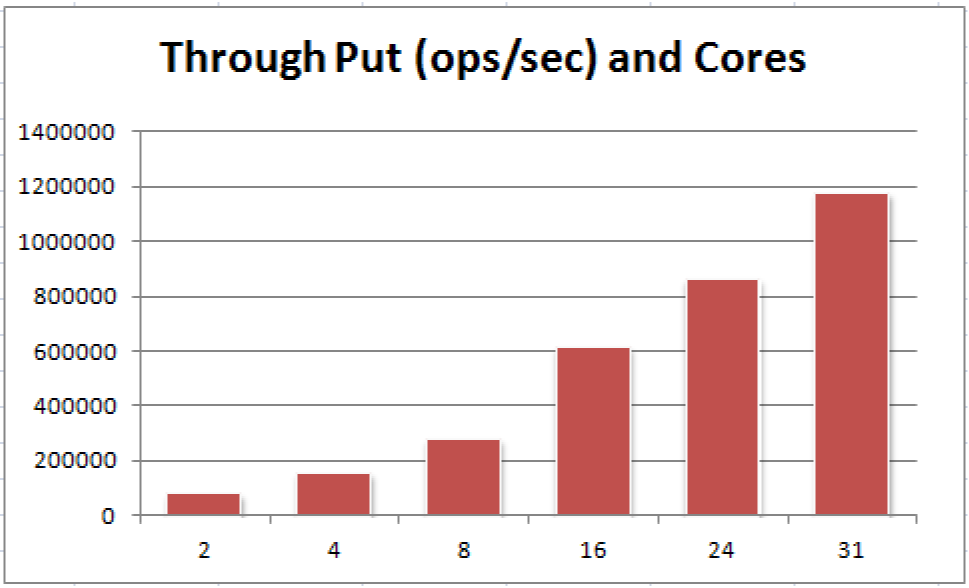

Figure 1: Oracle Internet Directory Scalability

HARDWARE AND OS CONFIGURATION

OID/DB Hardware details

OID and database co-existed in the same hardware.

CPU	Oracle SPARC T7-2 (2 chip/ 64cores), 4.13GHZ Test Config: 1 chip (32 cores) virtualized with OVM for SPARC
RAM	512GB
IO	2 X Internal disks 2 X 10 Gigabit Ethernet card
Storage	ZS3-2
OS	Oracle Solaris 11.3

Client Node details

SLAMD tool was used to generate the LDIF file and run the benchmark test. We used 1 node T7 for slamd client.

CPU	Oracle SPARC T7-1
RAM	256GB
IO	2 X Internal disks 2 X 10 Gigabit Ethernet card
OS	Oracle Solaris 11.3

TEST CONFIGURATION

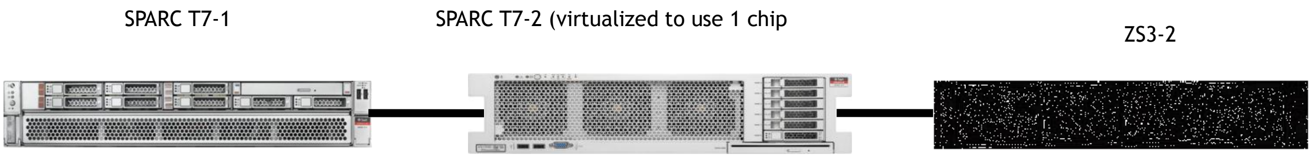

Figure 3. Oracle Internet Directory Deployment Architecture Overview

To run the benchmark we used the SLAMD distributed load generation engine. SLAMD clients were deployed on a T7-1 node (Application Tier) and connected to the OID server instances on a single SPARC T7-2 LDOM with 32 cores enabled. The backend OID database was run on the same SPARC T7-2 LDOM (Data Tier). All systems were inter-connected via 10GB network.

INSTALLATION

The installation sequence is as follows

DATABASE

We installed database 12c (12.1.0.2) on SPARC T7

OID

We installed OID 11gR1-PS7(11.1.1.7.11) on SPARC T7.

DATABASE TUNING

Following tuning was performed for database-

Parameter	Value	Description
Sga_max_size	96G	The hard limit up to which sga_target_size can dynamically adjust sizes for the System Global Area (SGA).
Sga_target	96G	The target size for the System Global Area
Pga_aggregate_target	16G	Target size for the Program Global Area.
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	
Audit_trail	None	
Processes	5000	Maximum number of processes

DATABASE Redo Logs

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

OID TUNING

We kept most of the configuration in OID as OOTB; typically we change process and thread configurations for every deployment based on the hardware configuration. Following parameters were changed for this benchmark.

Parameter	Value	Description
Orclserverprocs	32	Number of server process.
Orclmaxcc	8	Number of worker threads per server process. This was the optimal configuration for this hardware.
Orclecacheenabled	2	Enable entry cache & Result set cache.
orclecachemaxsize	260g	260G needed to cache 50 million users.
Orclskiprefinsql	1	Skip referral processing in SQL, since there are no referral entries in this test.
Orclextconfflag	0	Turn off audit and Ecid generation.
Orclmatchdnenabled	0	Disabled match DN, this does not affect successful search.
orclecachemaxentries	50000000	50 million entries to cache.
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

OID ACL

All tests were run with new attribute level Access Controls defined at the root node of each search base DN to explicitly grant attribute level access to the attributes that were requested in the search. Note that in OID, the admin user also comes under the purview of Access Control i.e. one could define ACL's to deny access to the admin user as well. Hence, in addition to these new ACL's, there are several out-of-box ACL's that come into play by default.

BULKLOAD

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 that 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

OID Database Size After the Bulkload

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

THE BENCHMARK TEST SCENARIOS

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 and UID are the required attribute
- Each LDAP search operation matches a single entry
- The total number concurrent clients was 1000 and were distributed amongst 2 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 clients lookup the same entry and all entries are searched **randomly**.
- In one run of the test, random entries from the 50 Million entries are looked up in as many LDAP Search operations.
- Test job was run for 60 minutes.

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=", 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.

LDAP Compare Operations Test

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

- SLAMD 'CompareRate' job was used.
- No client compared user password of the same entry i.e. only unique entry was used by each client
- BaseDN of the Search is root of the DIT.
- Each LDAP Compare operation matches user password of user.
- The total number concurrent clients was 1000 and were distributed amongst 2 client nodes
- Each client binds to OID once and performs repeated LDAP compare operations.
- In one run of the test, random entries from the 50 Million entries are compared in as many LDAP compare operations.
- Test job was run for 60 minutes.

LDAP Modify Operations Test

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

- SLAMD 'LDAP ModRate' job was used.
- A total of 50 concurrent LDAP clients were used.
- Each client updates a unique entry each time and a total of 50 Million entries are updated
- Test job was run for 60 minutes.
- Value Length was set to 11.
- Attribute that is being modified is not indexed.

LDAP Mixed Load Test

The test scenario involved both the LDAP Search and LDAP Modify clients enumerated above.

- The ratio involved 60% LDAP Search clients, 30% LDAPBind and 10% 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 60 minutes.

LDAP Add Load Test

The test scenario involved concurrent clients adding new entries as follows.

Slamd standard Add Rate job is used.

- 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).

THE BENCHMARK TEST 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	1177947	0.847
CompareRate	1000	650047	1.537
AuthRate	1000	414394	2.41
Modify	50	22408	2.23
Mix **BIND** **MODIFY** **SEARCH**	1000 300 50 650	320071 96310 16052 208658	3.12 3.57 7.67 2.56
Add	16	1436	11.094

CPU SCALABILITY TEST (SPARC T7 AND T5)

Number of CPU Cores	Number Of OID server process (orclserverprocs)	clients	Through put (ops/sec) T7	Latency (msec) T7	Through put (ops/sec) T5	Latency (msec) T5
2	2	100	79300	1.26	68399	1.46
4	4	100	156114	0.639	145879	0.68
8	8	500	280029	1.78	270601	1.84
16	16	500	615563	0.811	560709	0.88
24	24	1000	863343	1.156	823741	1.21
32	32	1000	1177947	0.847	944624	1.05

LDAP Search Operations (32 cores)

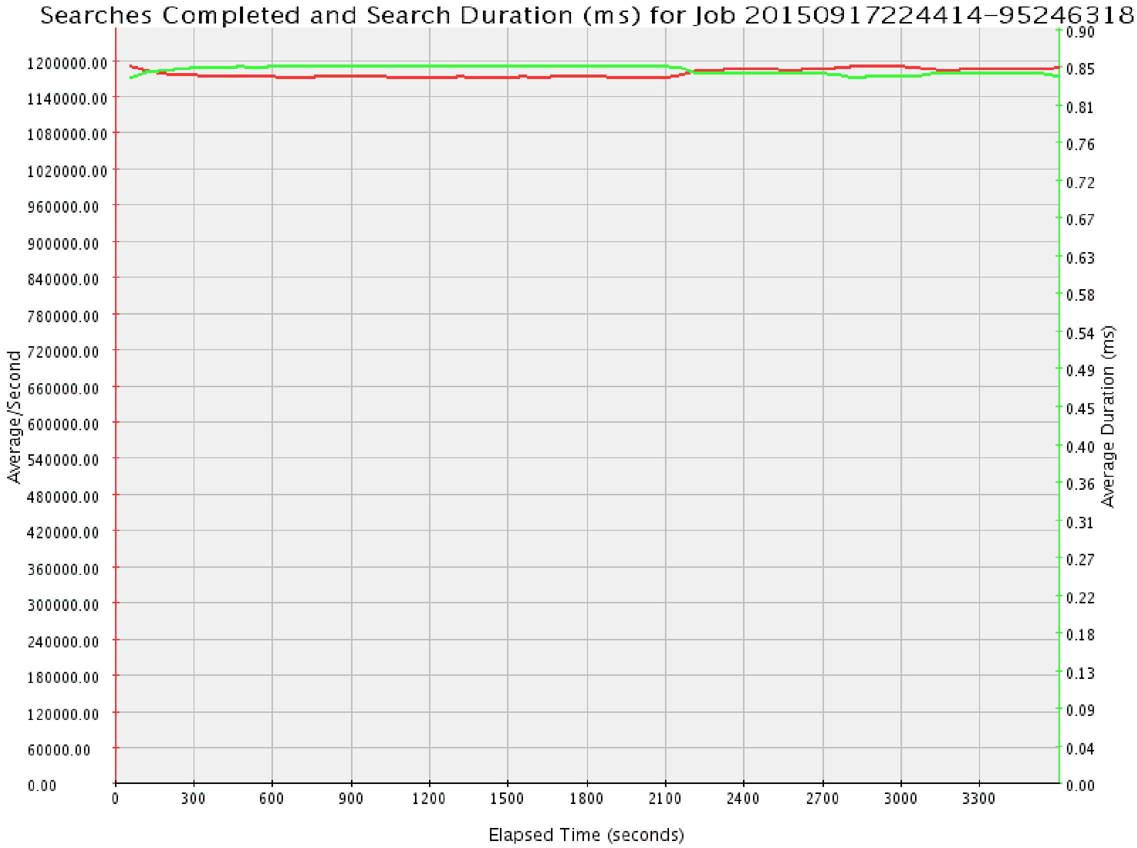

LDAP Compare Operations (32 cores)

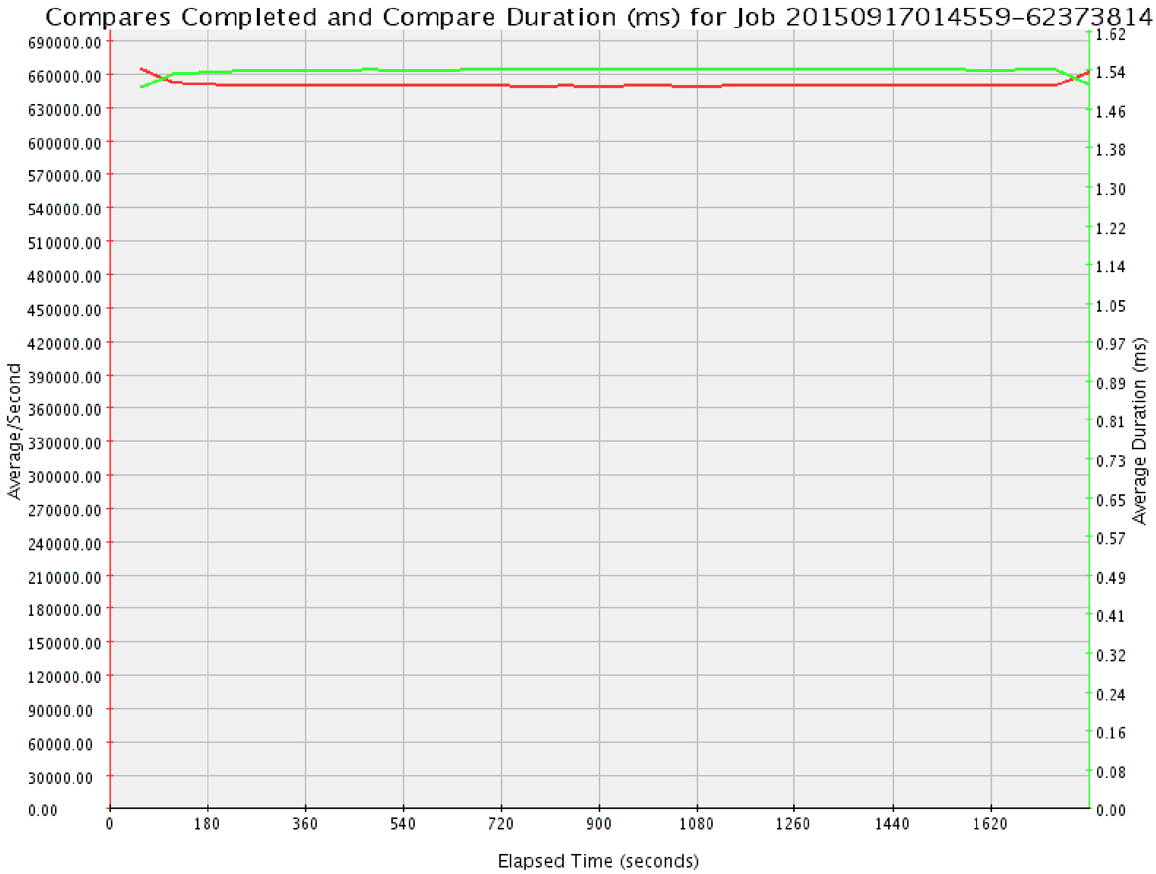

LDAP Authentication Operations (32 cores)

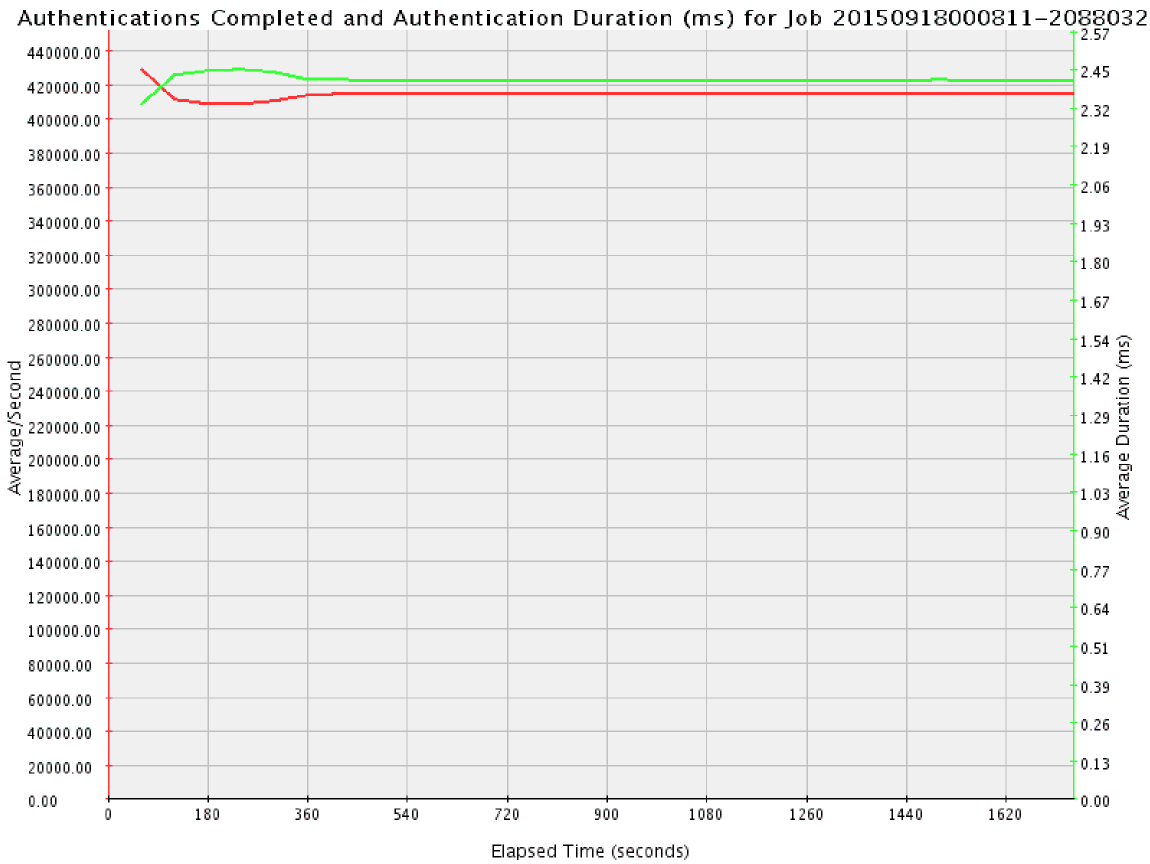

LDAP Modify Operations

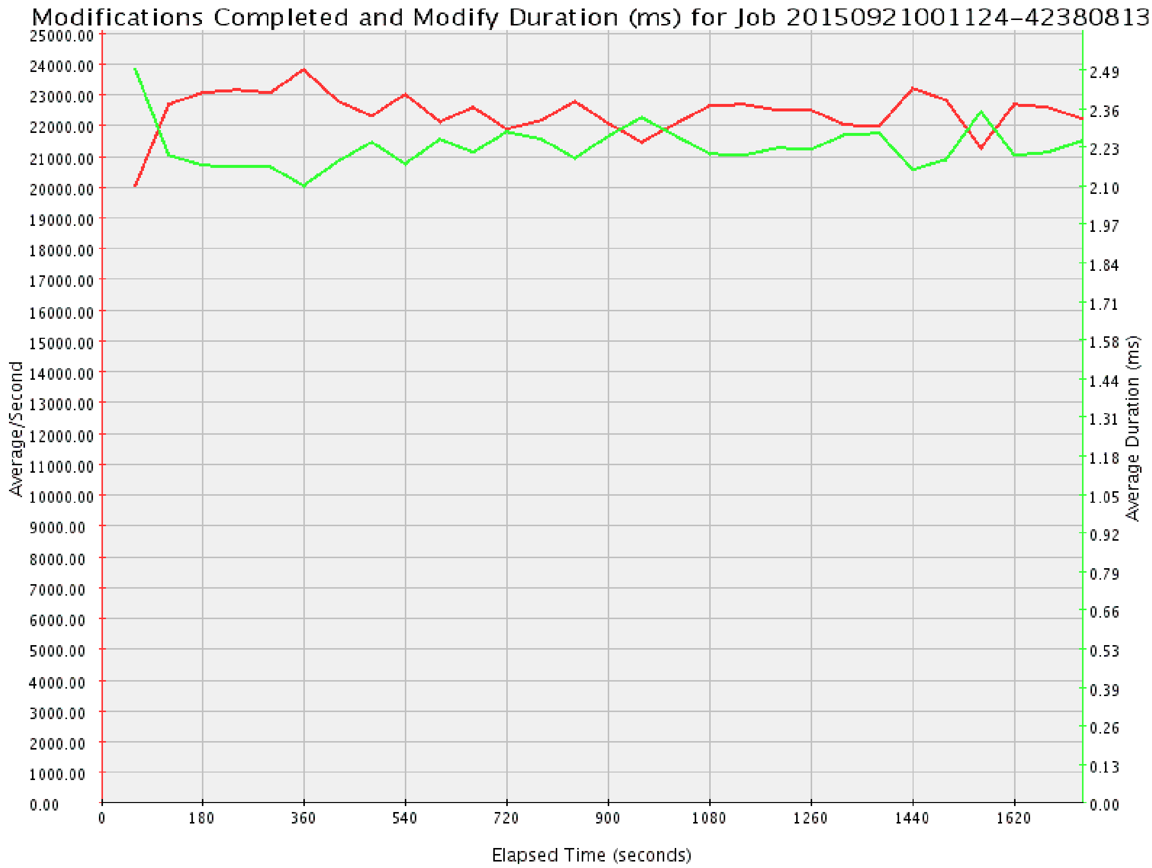

LDAP MIX Operations

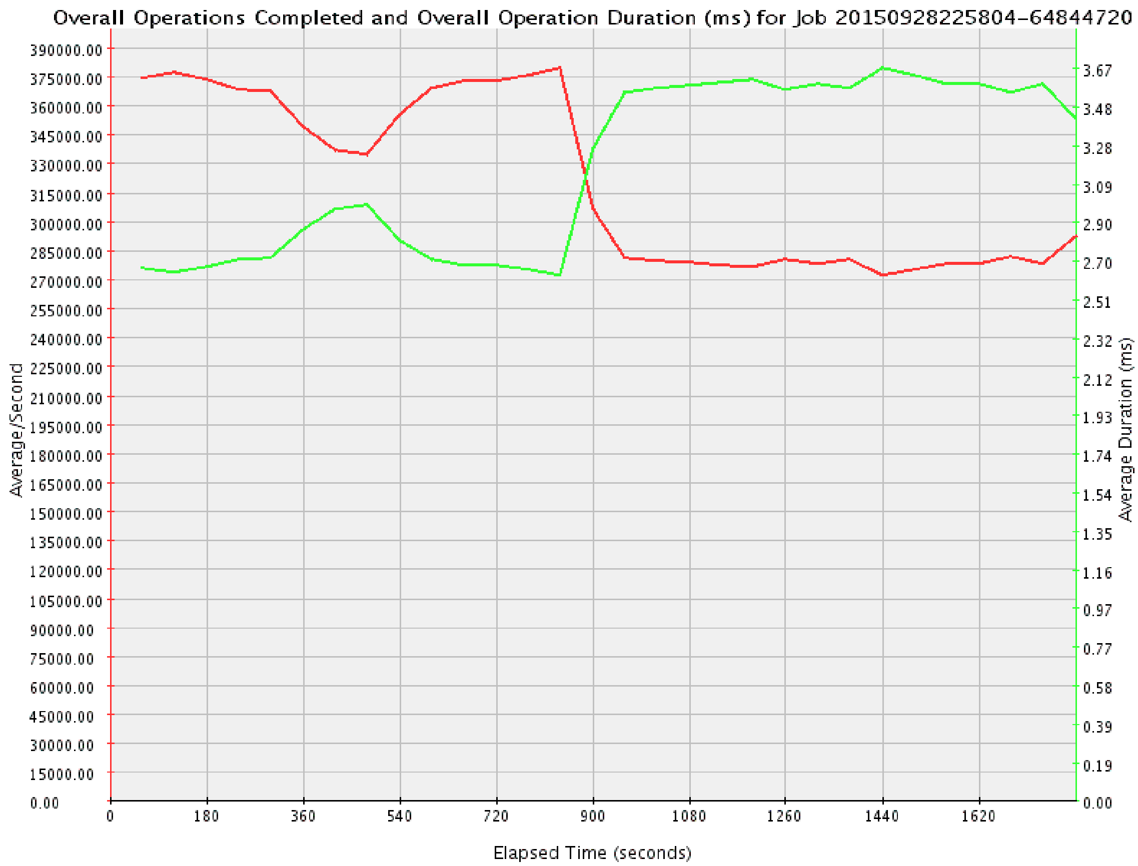

LDAP ADD Operation

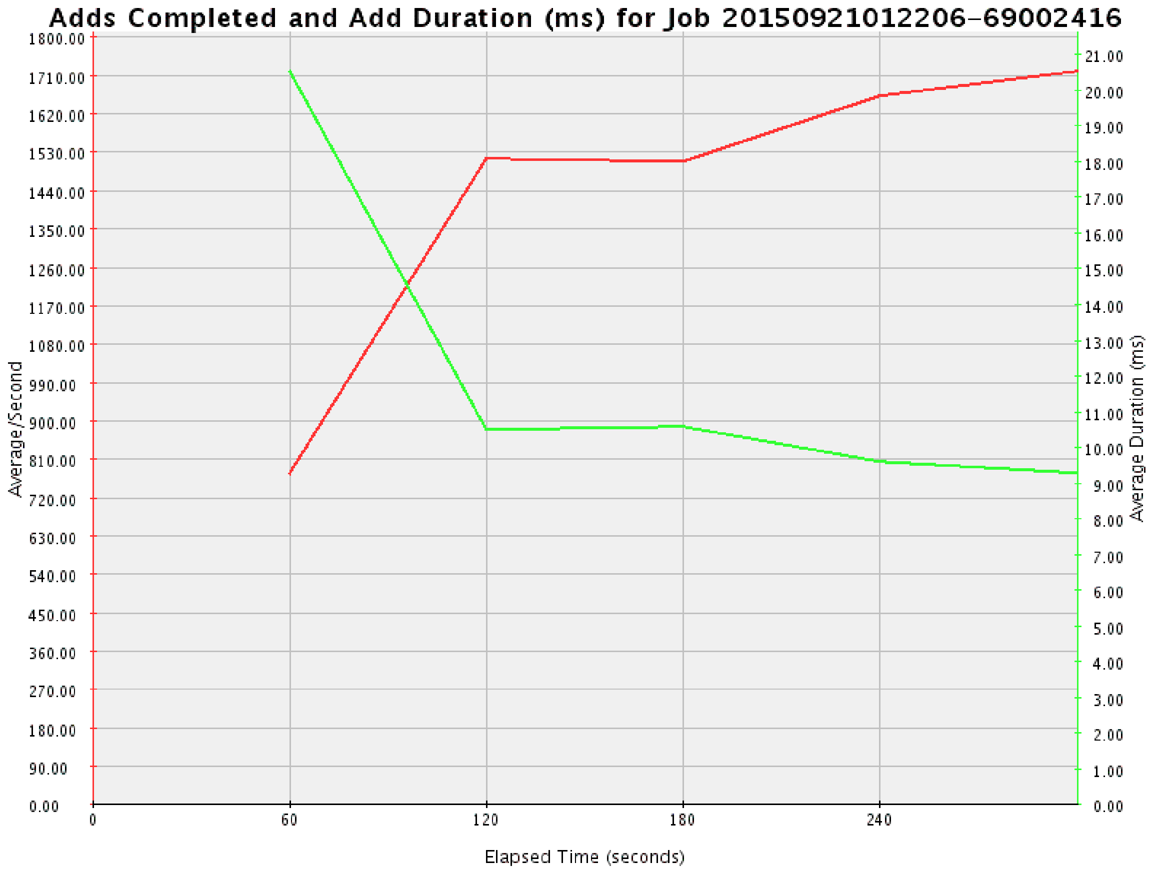

LDAP Search Operations (24 cores)

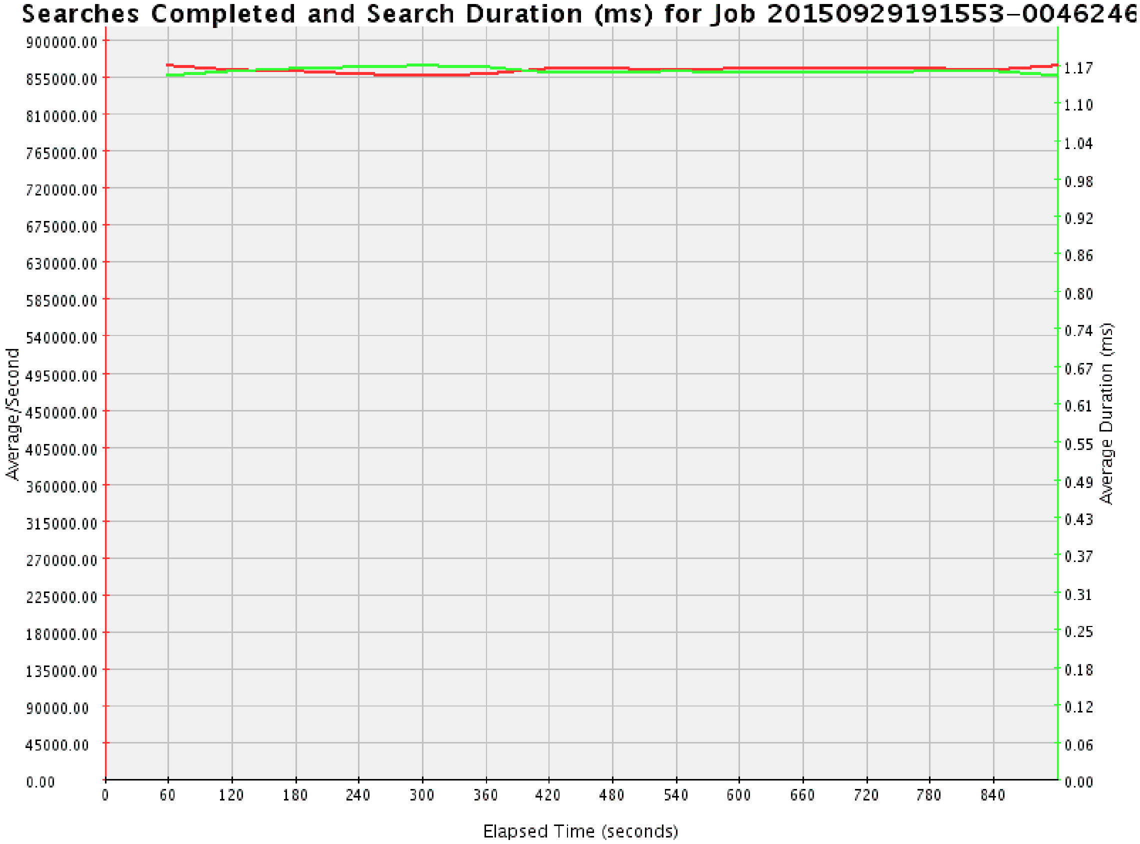

LDAP Search Operations (16 cores)

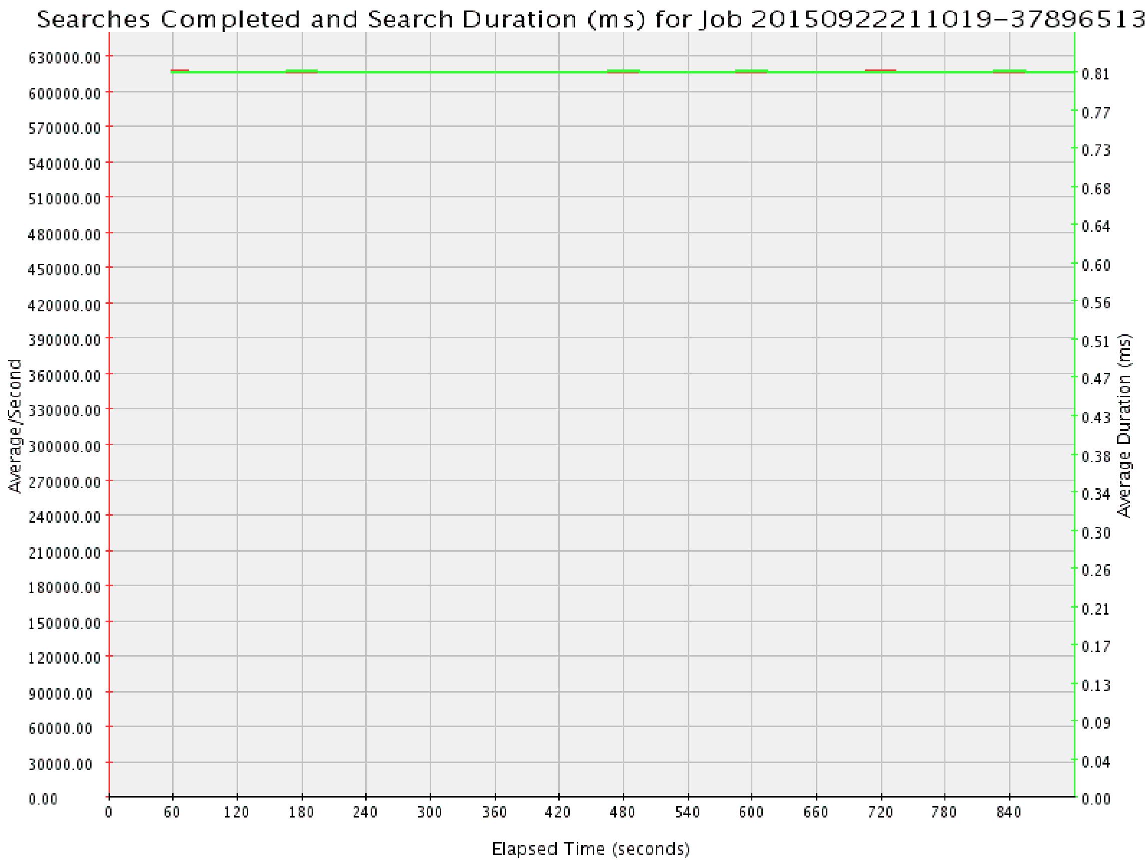

LDAP Search Operations (8 cores)

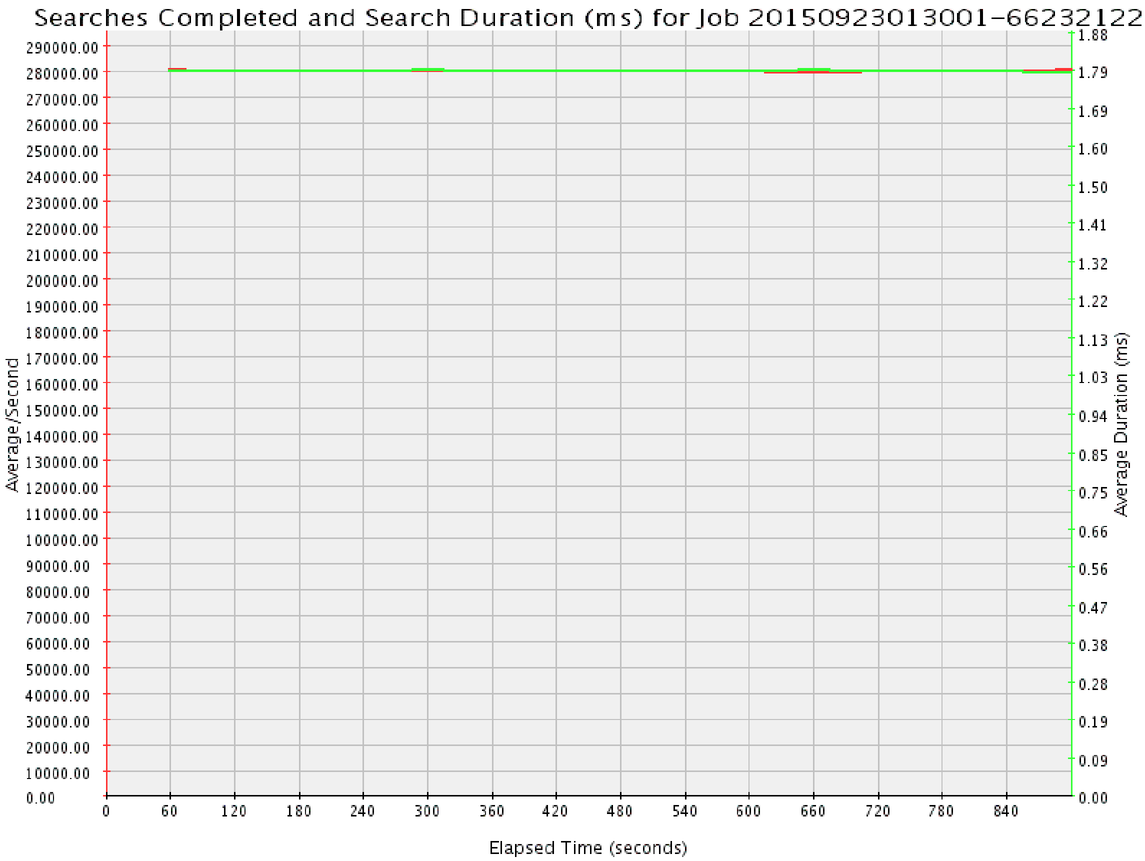

LDAP Search Operations (4 cores)

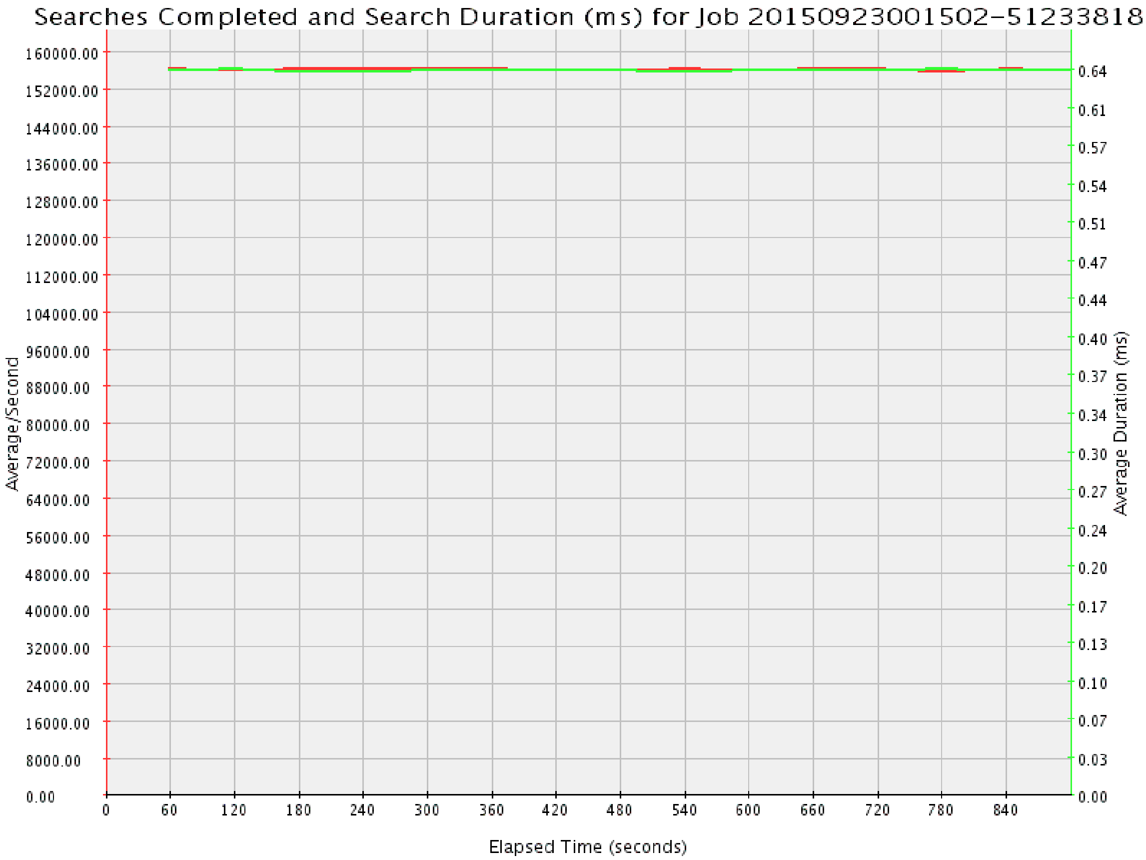

LDAP Search Operations (2 cores)

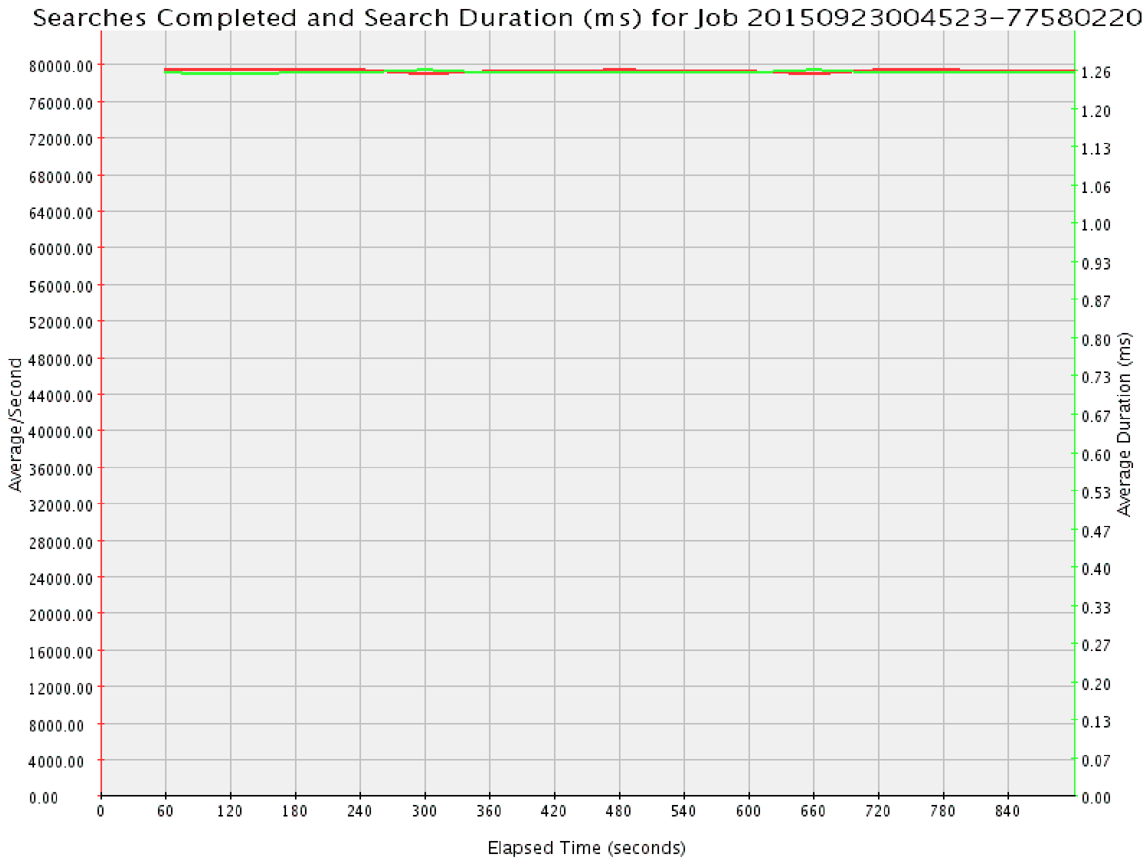

CONCLUSION

The following conclusions were drawn from the benchmark –

- OID search latency improved by more than 24% on SPARC T7 compared to SPARC T5
- OID LDAP Search throughput hits 1M on a single T7 chip, with linear scaling from 2 to 31 cores
- LDAP servers form the critical backbone for authentication and authorization services and OID is very suitable for such deployments given its excellent availability, scalability, and performance and manageability characteristics.
- OID on single T7 configured with 1 chip can deliver the performance needed by many Telco or cloud environments.

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/home/oid/app/ram/fast_recover
y_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.

Oracle Internet Directory, 50 Million User Directory benchmark on SPARC T7-2
October 2015

Author: Hanli Ren
Contributors: Dipankar Thakuria, Ramakrishna Bollu, Uday Shetty,

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