An Oracle White Paper
December 2011

Oracle Internet Directory 11g Benchmark With 10 Million Users On Oracle Exalogic X2-2

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) 11.1.1.6 and OID replication.

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 276,000 search operations / sec per X4170M2 node, with 10M user entries, Oracle Internet Directory showcases linear server performance resulting in 1.7 million operations / sec on six X4170M2 nodes. Not to mention there's no upper limited on performance scalability.

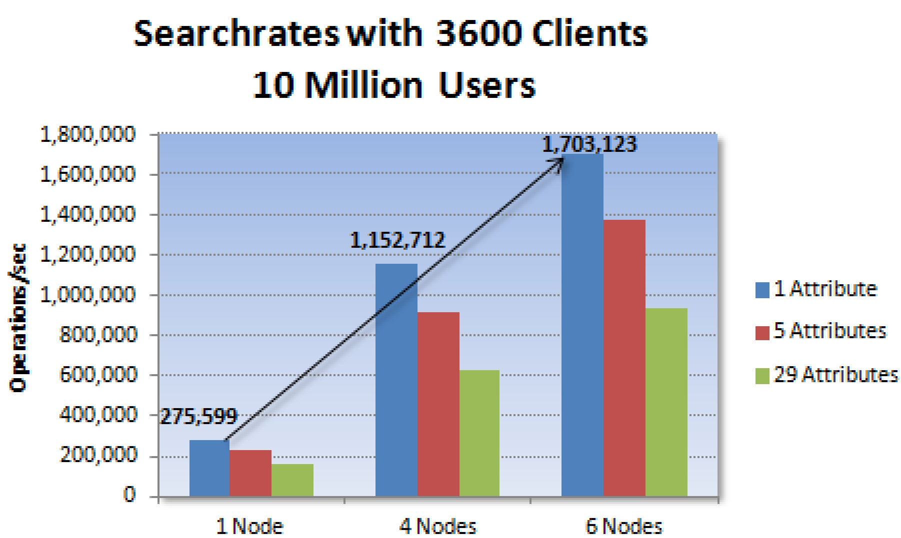

Figure 1: Oracle Internet Directory Scalability

This paper is the second one in a series of benchmarks performed on top of the same class of hardware (Oracle Sun Fire X4170), so the results can be put into perspective to each other.

The two other papers are:

1. "Oracle Internet Directory 11g and Oracle Exadata Database Machine in the Facebook age". Here the focus is on scalability and performance based on Oracle DB with 500M users leveraging Oracle Exadata Database Machine X2-2, and achieving 440.000 operations / sec, without using OID server caching or data partitioning.

2. "Oracle Internet Directory 11g scalable architecture with 1 million users". The focus is the unique server architecture (multi threaded, multi process), and the resulting scalability on one physical hardware node. Here we are comparing Oracle Internet Directory performance with OID server cache enabled / disabled.

Please see those papers for further details, or contact the authors of this whitepaper.

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

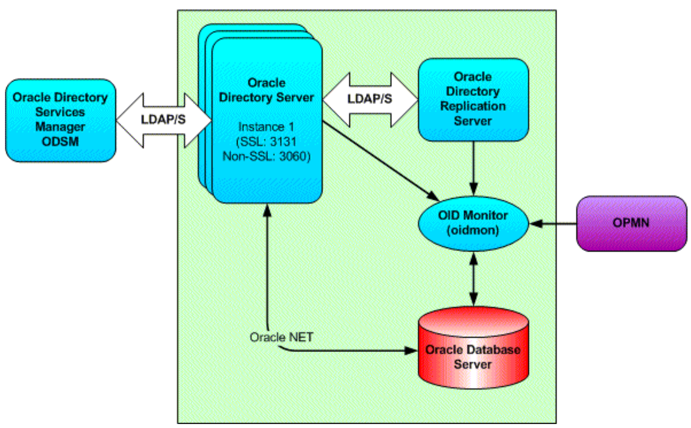

Figure 2: Oracle Internet Directory typical node architecture

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 this bechmark we made use of of OID's flexible deployment options to maximize performance with Oracle Exalogic hardware. The OID servers were deployed on 6 dedicated hosts (see Figure 5) to maximize host resource utilization.

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 Exalogic X2-2 Overview

The Oracle Exalogic Elastic Cloud is an Engineered System, consisting of software, firmware and hardware, on which enterprises may deploy Oracle business applications, Oracle Fusion Middleware or software products provided by Oracle partners. Exalogic is designed to meet the highest standards of reliability, serviceability and performance under widely varied, performance-sensitive, mission-critical workloads.

Exalogic dramatically improves the performance of virtually any standard Linux, Solaris and Java application1 with no code changes required and reduces application implementation and ongoing costs versus traditional enterprise application platforms and private clouds assembled from separately sourced components provided by multiple competing vendors.

In real-world performance testing and production deployments Exalogic typically outperforms other platforms by wide margins, often delivering two to ten (or more) times the application performance. Not only do applications become more responsive and deliver a greatly improved user experience, they are also more resource efficient.

With an Exalogic system, enterprises are able to support any given application workload with less hardware, less power, less heat, less data center space and less software. Because the Exalogic system is fully pre-integrated by Oracle it is also easier to provision, manage and maintain, further reducing ongoing costs and shortening time to value for new projects. Exalogic systems are designed for high availability and zero-down-time maintenance and can be scaled linearly from a single, Quarter Rack configuration to a large system of eight Full Racks with no service disruption and no additional external hardware required.

	Quarter Rack	Half Rack	Full Rack	2 - 8 Racks
2.93 GHz Xeon Cores	96	192	360	720 - 2880
1333 MHz RAM	768 GB	1.5 TB	2.8 TB	5.6 – 22.4 TB
FlashFire SSD	256 GB	512 GB	960 GB	1.9 – 7.7 TB
SAS Disk Storage	40TB	40TB	40TB	80 – 320 TB
All figures are model EL X2-2	8 Compute Nodes	16 Compute Nodes	30 Compute Nodes	

Figure 4 - Exalogic Configurations

In our benchmark we used a Quarter Rack each for the OID servers and the SLAMD clients for load generation.

[1] Any application that supports Oracle Linux (version 5 update 5 or later, Unbreakable Enterprise Kernel, 64-bit), Oracle Solaris 11 Express (or Solaris 10 Zone, x86)

Benchmark Deployment Architecture

Distributed Architecture

As outlined in the Oracle Internet Directory introduction we used a distributed three tier architecture to simulate a real world deployment. From a high level perspective the setup is outlined in Figure 3.

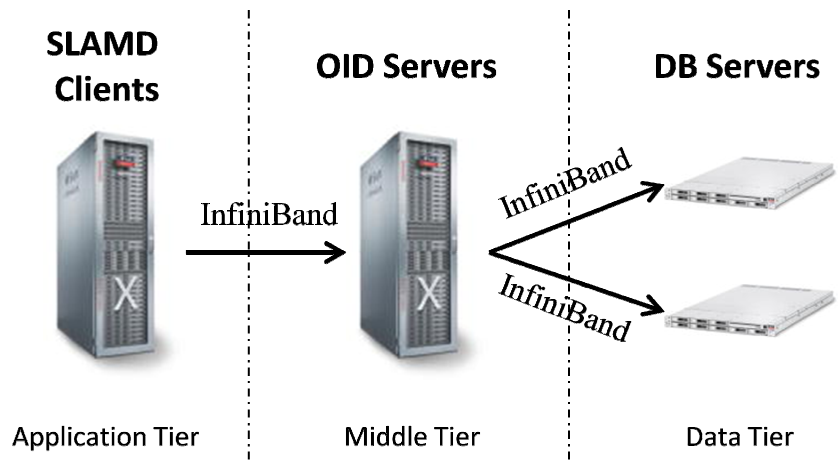

Figure 3. Oracle Internet Directroy Distributed Deployment Architecture Overview

To run the benchmark we used the SLAMD distributed load generation engine. SLAMD clients were distributed across seven nodes of an Exalogic quarter rack (Application Tier), and connected to the SLAMD server on the eighth node. OID server instances were installed on an Exalogic quarter rack (Middle Tier). The SLAMD server by itself distributed the load across six Oracle Internet Directory Servers. The backend OID database was run on two Sun Fire X4170 (Data Tier).

All systems were inter-connected via Infiniband network.

The actual relationship between software components and hardware nodes is outlined below in Figure 5.

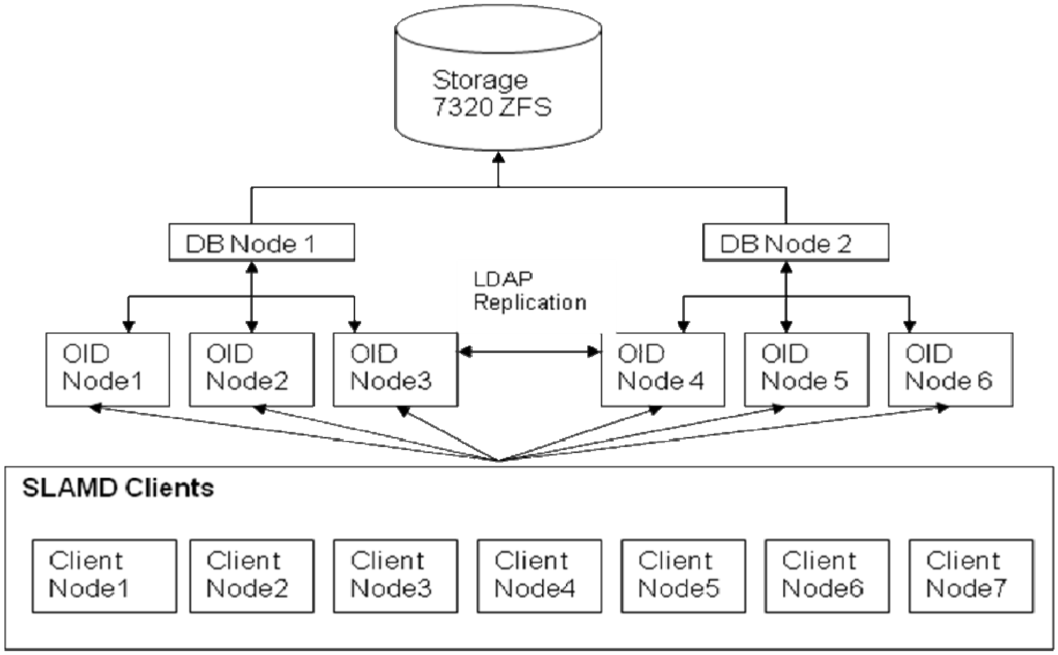

Figure 5. Oracle Internet Directory Distributed Deployment Architecture Details

Database Tuning

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

Parameter	Value	Description
SGA_MAX_SIZE	24GB	The hard limit up to which sga_target_size can dynamically adjust sizes for the System Global Area (SGA).
SGA_TARGET	24GB	The target size for the System Global Area.
PGA_AGGREGATE_TARGET	718MB	Target size for the Program Global Area.
audit_trail	none	Configures database auditing.
Processes	1500	Number of database processes.
open_cursors	300	Maximum number of cursors a database session can have.

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	12	One server process per CPU core. Since there were 12 cores on system, we keep one process per core.
Orclmaxcc[1]	4	Number of worker threads per server process. This was the optimal configuration for this hardware.
Orclecachemaxsize	66GB	Maximum size of the entry cache. This was based on the available memory in the hardware.
Orclecachemaxentries	10 Million	Maximum number of entries that can be cached. This was calculated based on entry size and cache size.
Orclextconfflag	0	Disabled auditing logs.
Orclstatsflag	0	Disabled collection of statistics information.
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.

Data Characteristics

Oracle Internet Directory (LDAP Data)

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

- 38 attributes (12 generated operational attributes), of which
  - 1 was a binary attribute with length 1.5kB
  - The others had cumulative length of 1.5kB
- Total entry size was 3kB

This was chosen to represent a real-world deployment where the user entry may contain binary information such as a photograph or cryptographic identifiers like certificates.

In addition, 10 group entries were created with the following properties:

- 1 million members each
- Each user is a member of one group
- Each entry was 450MB in size

This was chosen to represent the high-end grouping requirements in a deployment of this size.

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

Oracle Internet Directory Database

The database representation of this data had the following characteristics:

TableSpace Name	Size in GB
OLTS_ATTRSTORE	26
OLTS_CT_STORE	34
OLTS_BATTRSTORE	20
SYSAUX	1.3
OLTS_DEFAULT	0.001
OLTS_SVRMGSTORE	0.0002

Workload Scenario

The workload scenario tested was a 10 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 10 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 3600 and were distributed amongst 6 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**.
- Test job was run for 60 minutes.

Additionally this test was run with 250 concurrent clients against 1 client node to determine low latency metrics. Furthermore, this test was run with 5 required attributes and all attributes to determine performance for searches obtaining more data.

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 3600 and were distributed amongst 6 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.

Additionally this test was run with 250 concurrent clients against 1 client node to determine low latency metrics.

LDAP Group Search Operations Test

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

- SLAMD ‘CompRate’ job was used.
- The total number of concurrent clients was 3600 and were distributed amongst 6 client nodes
- Each client does subtree search with the filter member=’RandomDN’.
- Each lookup will return one group entry.
- Required attribute is set to DN.
- Test job was run for 60 minutes.
- Search performed such that all the 10 group entries were touched randomly.

Additionally this test was run with 3600 concurrent clients against 1 client node to determine low latency metrics.

Benchmark Results

Exhaustive Workload targeting 10 million entries

High concurrency scenario

Number of Clients: 3600

Number of OID Nodes: 6

Test Scenario	Through put (ops/sec)	Latency (msec)
Search (Random) 1 attribute	1,703,123	2.111
Search (Random) 5 attributes	1,377,699	2.610
Search (Random) all attributes	936,835	3.839
Authentication (Random)	648,113	5.551
Group Search (Random)	234,016	3.389

Low latency scenario

Number of Clients: 250

Number of OID Nodes: 1

Test Scenario	Through put (ops/sec)	Latency (msec)
Search (Random) 1 attribute	275,599	0.904
Search (Random) 5 attributes	229,725	1.085
Search (Random) all attributes	158,900	1.570
Authentication (Random)	102,356	2.439

Results Analysis

Cache

In OID 11.1.1.6.0 release entry cache is enhanced to support distributed architectures. This benchmark leverages this entry cache to provide high throughput. With 66GB of cache memory it was feasible to have most of the user entries (10 million) in the cache. As a result the vast majority of test operations exercised in this workload benefited from significant cache hits.

Since previous versions of OID did not support a shared cache, multi-instance deployments never benefited from cache performance. With OID’s support for a shared cache across multiple server instances on the same node, now customers can also benefit from cache performance.

This opens up a host of deployment options for customers. For example, a customer can have multiple OID nodes with each running multiple instances of OID to maximize hardware utilization. As standard in an OID deployment, each OID node will be looking at the same database. However, while previously every query results in a hit against the DB, now each OID node will build up a cache based on the traffic it sees. This means that over time, each node will have cache contents optimized to service the traffic it sees. Additionally, load on the DB will decrease as an increasing number of requests are served out of the cache and as a result increase throughput on the OID nodes by as much as 5x.

Data Partitioning

The workload scenario tested here did not require data partitioning. The entire set of 10 million user entries was stored in a single Directory Information Tree (DIT) without compromising performance. Consequently, no a priori knowledge of the expected workload characteristics was required to partition for maximum performance. Furthermore, no subsequent maintenance, such as repartitioning, was required to address changes in workload characteristics over time. This illustrates how OID's ability to handle all data within a single DIT empowers administrators to load-and-forget their data thus reducing administrative overhead.

It is important to note here that the Cache enhancements done above also complements OID’s paradigm of scaling without data partitioning. When bringing in new LDAP client applications on board in a data center and the need to add capacity arises, one can simply add additional OID nodes to service that traffic without worrying about partitioning data or having an understanding of workload characteristics. As stated above, as that node is used for the application it will optimize it’s cache to meet the needs of the request load it services without administrator intervention. This significantly reduces administrator burden during scale-up and load balancing exercises in the data center.

Large Entry Size

This benchmark operated on user entries that were 3kB in size (including a 1.5KB binary attribute). Nevertheless, since OID's architecture allows it to operate on each individual element of an entry, no discernible performance impact was observed.

Large Group Size

Part of this exercise also involved operations on groups with 10 million members. To put things in perspective, our entire workload scenario would be covered by the members of a single such group. Nonetheless, OID was able to service these requests without hindrance at high throughput.

Latency

The results confirm the intuitive notion that latency increases with concurrency (ie. the number of concurrent clients). As the number of clients connecting with OID grows, the amount of resources used and context switches that occur within the OID server also grows. It is clear that this would cause an increase in latency. While it was feasible to reduce the number of clients connecting to OID in order to obtain better latency numbers, we believe that for a deployment of this magnitude at least 250 concurrent clients would be expected.

Concurrency

In order to showcase OID's ability to handle large numbers of concurrent clients, we ran the Random LDAP Search Operations Test on the workload scenario for 250 and 3600 clients respectively. operation latency increased as expected.

Sustained Uniform Performance

As can be seen in the performance graphs for our benchmark tests in Appendix C & D, OID was capable of maintaining uniform performance levels with no spikes throughout each test. This showcases OID's ability to maintain high performance numbers over long periods of time without service disruption.

Conclusion

The new Oracle Internet Directory 11.1.1.6 entry cache provides a performance boost of up to 300% above the previous implementation. Together with Oracle's Internet Directory scalable and highly available architecture it provides a robust foundation for very large directory deployments that require outstanding and sustainable performance.

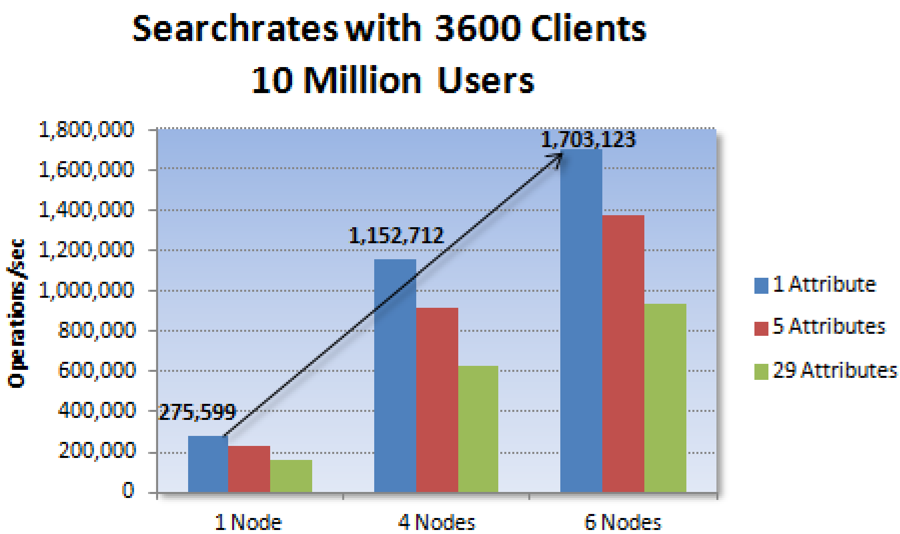

Figure 6: Oracle Internet Directory Scalability

Appendix A: SLAMD Performance Graphs

High connections scenario
Exhaustive workload with 10 million entries
Number of clients: 3600, Number of nodes: 6

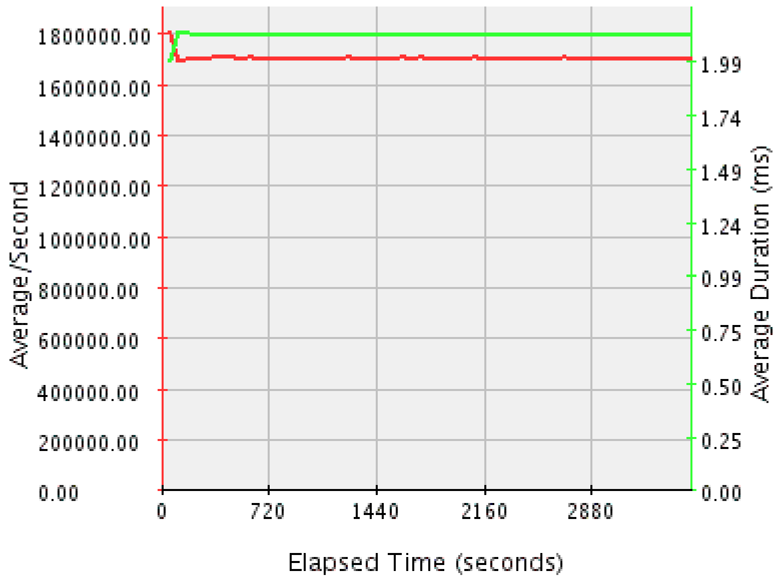

Search Rate with 1 Attribute

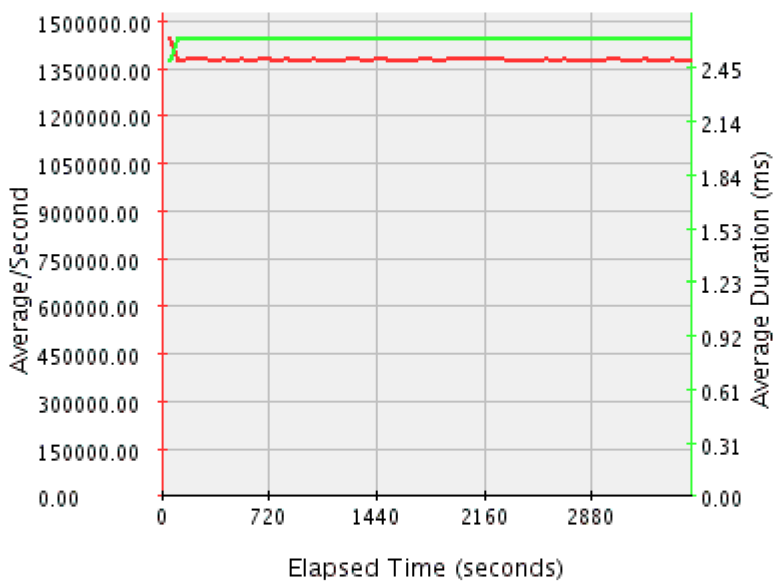

Search Rate with 5 Attributes

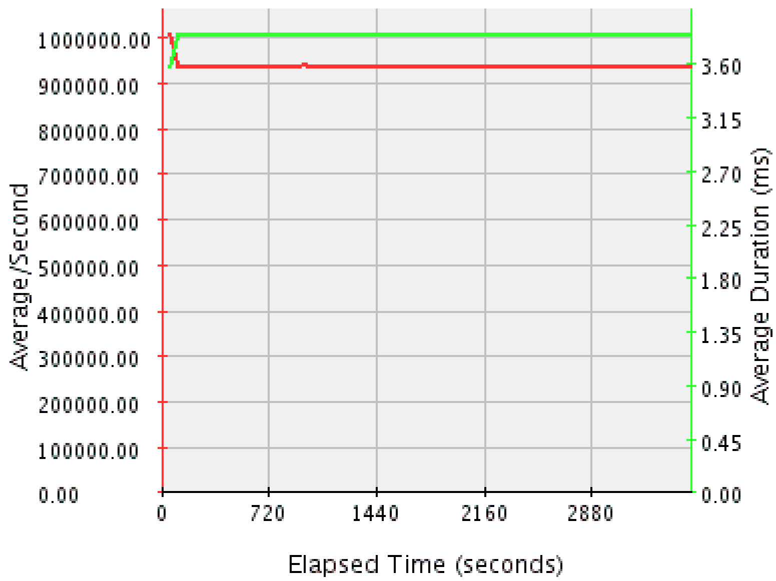

Search Rate with All Attributes

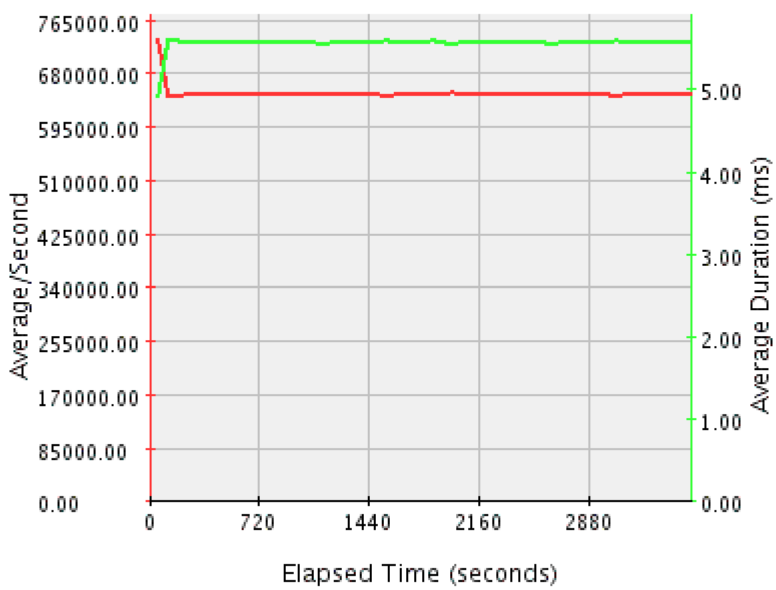

Authentication Rate

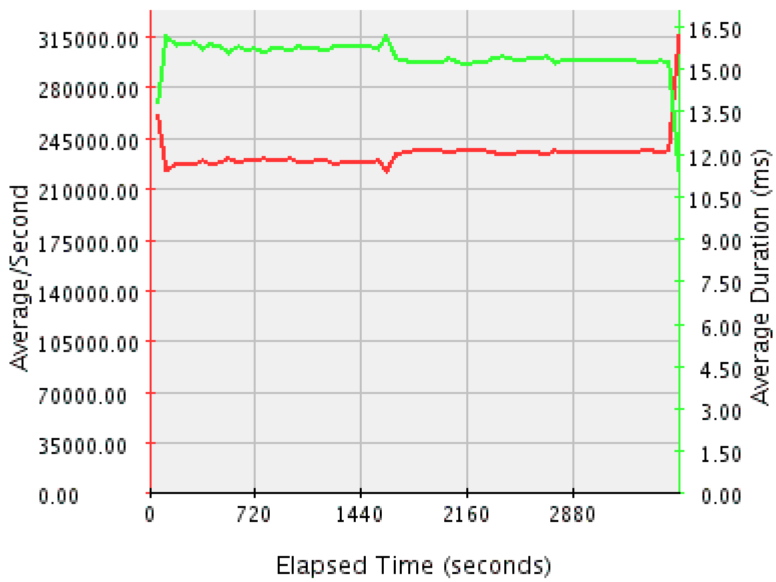

Group Search Rate

Low latency scenario

Exhaustive workload with 10 million entries
Number of clients: 250, Number of nodes: 1

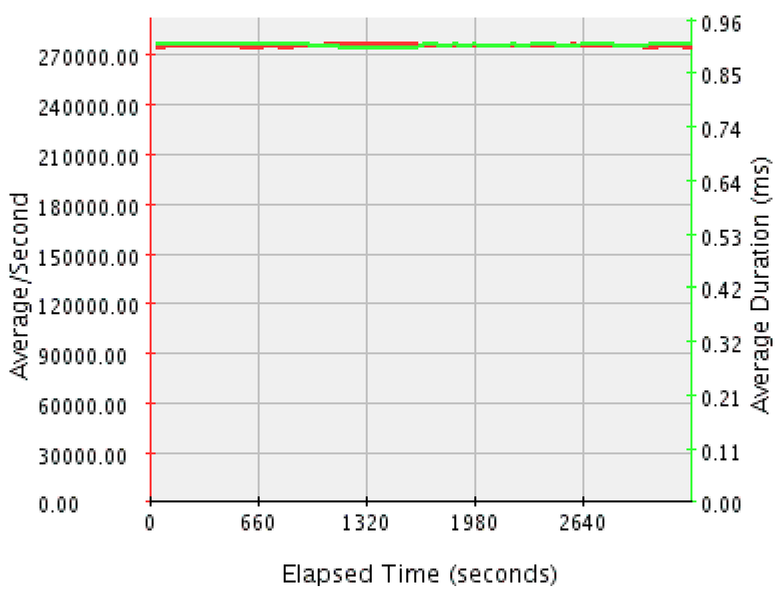

Search Rate with 1 Attribute

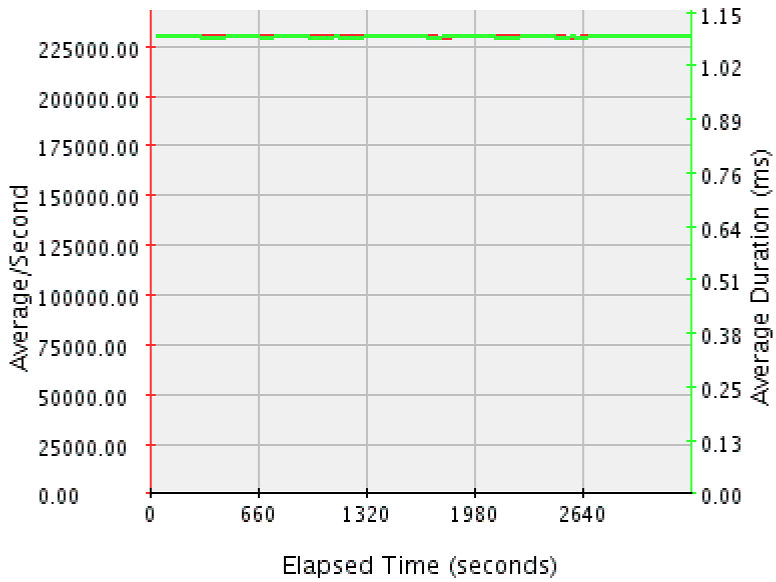

Search Rate with 5 Attributes

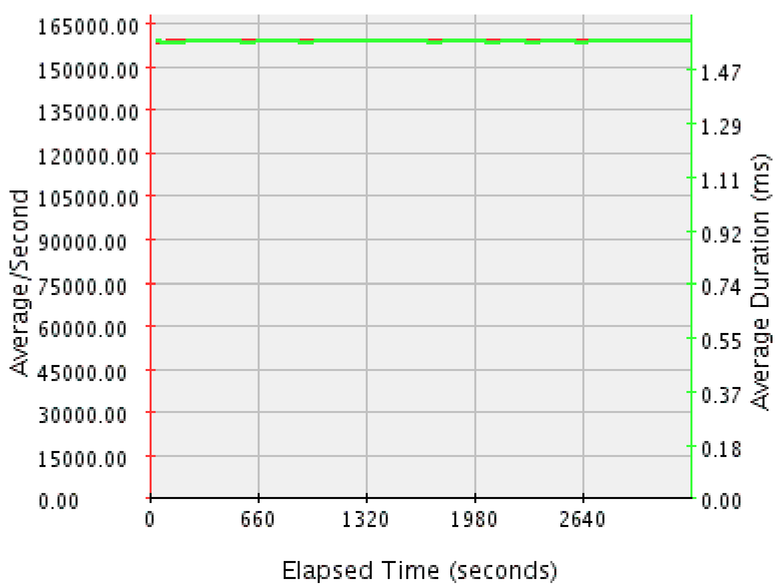

Search Rate with All Attributes

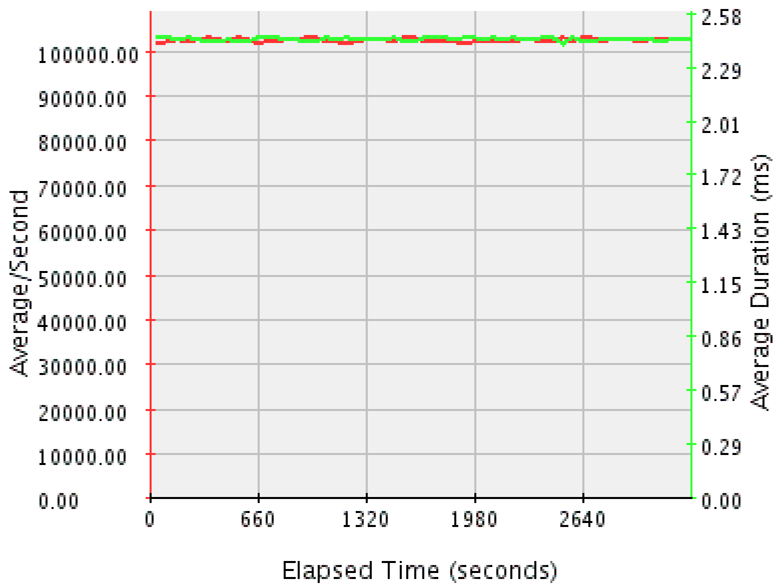

Authentication Rate

Appendix B: Template LDAP User Entry

SLAMD template file used to generate data:

```
define suffix=dc=us,dc=oracle,dc=com
define maildomain=oracle.com
define numusers=10000000

branch: [suffix]

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: welcome123
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}.
photo:: <base64:<random:alphanumeric:1500>>
```

Appendix C: Additional Information

1. Oracle Internet Directory 11g And Oracle Exadata In The Facebook Age
2. 2 Billion User Benchmark (Oracle Internet Directory 10.1.4.0.1)
3. Oracle® Fusion Middleware Enterprise Deployment Guide for Oracle Identity Management 11g Release 1 (11.1.1)
4. Oracle® Fusion Middleware High Availability Guide 11g Release 1 (11.1.1)
5. Oracle Internet Directory on the Oracle Technology Network
6. Oracle Exalogic Database Machine on Oracle Technology Network
7. Oracle Sun Fire 4170 Server
8. SLAMD Load Generation Engine

Oracle Internet Directory in the Facebook Age
December 2011
Author: Ramaprakash Sathyanarayan, Arun Theebaprakasam, Buddhika Kottahachchi
Contributing Authors: Olaf.Stullich@oracle.com

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