Directory Benchmark Comparison

ICSynergy
5601 Democracy Drive, Suite 205
Plano Texas 75024

Issued: November 29th, 2017

Sponsored by CA Technologies

Overview

There is a growing demand for business to have a tighter communication and collaboration with external parties such as business partners and customers. As a backbone for these services, Authentication Authorities and Directory services are key technologies for that evolution. They enable organizations to manage access both from-and-to external systems, including cloud services, in a consistent way.

Directory services are the core to Identity and Access Management and Federation technologies making them a key service to all organizations. While the vendor landscape is well established and relatively mature, the Directory service features, management, operations, deployment and capabilities vary among all vendors. This document focuses on capacity and performance, which are vital to the capacity, performance, scalability and reliability of the Identity and Access Management solutions used by organizations to interface with their consumers and partners.

This ICSynergy “LDAP Benchmark Comparison” documents a performance overview of the leading vendors in Directory services market segment.

Product and Versions Covered in this Comparison

Four directory services market leading vendors were chosen for the purpose of this benchmark:

- **Oracle**
- **Ping**
- **ForgeRock**
- **CA Technologies**

Oracle has two very different directory services products, ODSEE (Oracle Directory Server Enterprise Edition) which is in its end-of-life phase and it is now superseded by the Oracle Unified Directory (OUD). Because there are a significant number of large ODSEE deployments still in the market, we have included it in this comparison. The directory products from both ForgeRock and Ping are based on the Sun Microsystems open source LDAP server. The Ping Directory is the result of Ping's merger with UnboundID and has benefited from significant investment in engineering and development which distances it significantly from ForgeRock's OpenDJ despite having a similar heritage. CA Technologies acquired OpenDirectory and rebranded it initially as eTrust Directory in 2000 and CA Directory in 2006. CA Technologies directory service, is designed to be a high performance, easy to deploy and operate directory service with high levels of data integrity.

These five directory service products together represent the largest percentage of commercial LDAP licenses sold and deployed in the market.

The versions used in this comparison were the latest at the time of the testing/publishing of this document and are as follows:

Vendor	Product Name	Version
Oracle	Oracle Directory Enterprise Edition (ODSEE)	11g (11.1.1.7.2) 64Bit - Linux
Oracle	Oracle Unified Directory (OUD)	12cPS3 (12.2.1.3.0) - Linux
Ping	Ping Directory (Formally UnboundID)	6.0.1.0 GA 64Bit - Linux
ForgeRock	OpenDJ	5.5.0 64Bit GA - Linux
CA Technologies	CA Directory	12.6.01 (build 14046) Linux 64-Bit

Test Deployment and Architecture

There are a number of different testing models used by vendors, customers and market analysts when comparing software product performance or benchmarks. Software vendors have historically designed benchmarks around the understanding of operations on a per CPU core, as this provides a simplified set of metrics that can help customers in sizing and calculating capacity for a specific architecture. However, customers tend to test or compare software based on the specific design that they intend to deploy.

This comparison was based on a generic deployment design, (based on feedback and best practices from each vendor), to simulate a large, distributed directory deployment capable of supporting 100 million users. To an extent, this allowed the flexibility to simulate the effects an LDAP load has against a number of components instead of a single server instance. All the products were subjected to a similar architecture with the same LDAP load.

LDAP Proxy/Router

All the products vendors chosen with the exception of ForgeRock's OpenDJ, offer a Directory Component that handles LDAP specific load balancing operations. Each vendor may have a different name for their directory component (LDAP Proxy, Router, Virtual Directory etc), but the architectural benefits are similar across all the vendors. They allow the intelligent routing, rewriting, virtualization and load balancing of LDAP specific operations by inspecting the LDAP operation and forwarding the operation to the appropriate back-end server. LDAP Proxy/Routers are not directory repositories by nature and thus only handle logic capable of rewriting operations on the fly, or load balancing (or segregating) different types of operations (i.e. reads and writes).

For the purpose of the tests, we made use of LDAP proxies/routers only when recommended by the vendor and as part of their best practices with the exception of OpenDJ where no LDAP Proxy/Router was used.

AWS Instance Description

Each LDAP, Proxy or Load Server instance was configured with the exact specifications below:

- OS: Redhat Enterprise Linux 7.3 (HVM), SSD Volume Type
- Instance Type: m4.4xlarge
- Intel Xeon E5-2676 v3 2.4GHz
- vCPUs: 16
- Memory: 64 GB
- Storage Size: 150 GB, EBS Only
- Storage Volume Type: Provisioned IOPS SSD (IO1)
- EBS-Optimized: Yes
- Network Performance: High
- All traffic (load, management and replication) was routed through the internal AWS IP network to minimize network latency.

Test Protocol

The test was designed to generate a mixture of traffic load, that in percentage, would be similar to what a number of deployments sustain during a given peak hour for 100 million users. That is, a large number of authentication, attribute compare and profile search operations with a relatively small (9%-12%) amount of indexed attribute modification (in order to trigger the re-index of that attribute). The objective was to understand the latency of the response, and the percentage of active users that the infrastructure could sustain comfortably before any infrastructure upgrade was needed. We compared how the infrastructure behaves and reacts to this traffic load in terms of CPU utilization and how each product handles a given load. For this purpose, SLAMD was used, as an LDAP testing tool it is versatile, and built from the ground up to be an LDAP specific test bench. The results are a summary of a 30 minute SLAMD load run which consisted of four (4) *concurrent* traffic profiles to generate the load required for this test. Profiles used were:

- **AuthRate** - This test measured the Bind and Authentication Rate that the infrastructure is able to sustain within a period of time. Each operation includes a pooled bind and authentication for an application account that is used to search the test user ID's DN, unbind and then a bind attempt under that user ID to verify the password. This is a normal authentication flow and documents the number of authentications per second that a directory service is able to sustain. One AuthRate operation encapsulates the application bind/search of the DN and user binding into one operation.
- **CompareRate** - Compare Rate operations describe the cases when applications and services require the directory service to determine whether a specified entry has a particular attribute value. These operations are smaller than a search (as only a result code is returned instead of a set of attributes) and used often by services that need to handle transactions that perhaps do not require user authentication (i.e. a messaging infrastructure accepting an incoming email, but verifying that the user in-fact, exists). One CompareRate operation encapsulates the application pooled binding, searching, comparing and responding into one operation.
- **SearchRate** - Search Rate measures the rate at which an LDAP service can perform random, user-defined searches. The test included the bind/authentication of an admin account, the search/filter and retrieval of a given user profile and measuring the latency that the combined commands experienced on these operations. One SearchRate operation encapsulates the application pooled binding, searching, retrieving the object into one operation.
- **SearchMod** - Search Modification rate measures the rate at which an LDAP directory can perform a bind/authentication of an application admin account, perform a search of a user and then perform random, user-defined modifications to one of their attributes, simulating profile changes and triggering re-indexing at the LDAP server of a given attribute change. One SearchMod operation encapsulates the application pooled binding, searching and modifying an indexed attribute into one operation.

The rational for this mix is that during a peak hour, these are the operations that an infrastructure might experience with the most frequency, and measuring how many connections/operations each product is accepting and processing provides a sense of infrastructure capacity. While authentication and searches are the core of the operations that a directory service experiences, performing LDAP compares and modifications means that the product might have to balance the number of search operations that it is processing (or increase the latency to which the operation takes to complete). The traffic mix used for these tests were as follows:

	Number of Clients	Threads per Client	Percentage of Load Generated
AuthRate	14	4	40%
CompareRate	3	4	9%
SearchMod	4	4	11%
SearchRate	14	4	40%

Authentication/Search and Compare Architecture

When Authentication, Search and compare operations are a priority for the applications, all vendors recommend a flat design where a bank of directory servers are all concurrently sharing the LDAP requests to the applications (with a load balancer fronting the load.) The obvious downside of this design is a higher latency and fewer write operations as the information needs to be committed and replicated across a number of servers. Even though, this is the most favorable deployment model for authentication and search operations when supporting services that do not require a high level of writes as part of their design. All the products were subjected to a test that was comprised of 6 LDAP dedicated instances, to distribute load, to increase the effects of the replication agreements, to represent deployment that might have different geographical characteristics and to account for availability. As there are many possible models, a flat 6 server environment allowed for the combination of all other factors into one scenario.

The following diagram depicts the flat “authentication/search and compare” design used for acquiring the results of this test.

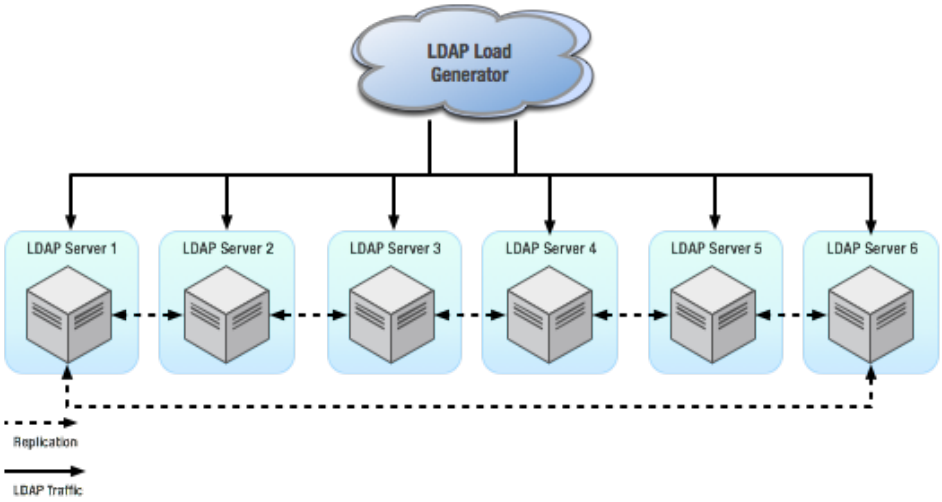

Modify/Write Architecture

The architecture and design of an LDAP service must take into consideration the expected volume traffic and the priorities given to the type of operations that the service needs to sustain. When the solution is expected to sustain a higher write/modifications/addition/deletion of objects and attributes the overall directory design may be modified to ensure enough capacity for these directory updates. For example, CA Technologies recommends a Horizontal Partitioned Configuration (HPC) as it segregates and minimizes the replication traffic needed and thus committing write operations at a higher volume when write operations are expected at a lower latency. CA Technologies does expect HPC to add some latency to the search/read operations as HPC prioritizes modification/write operations, but architectures where modification volumes are high, this is the recommended deployment model.

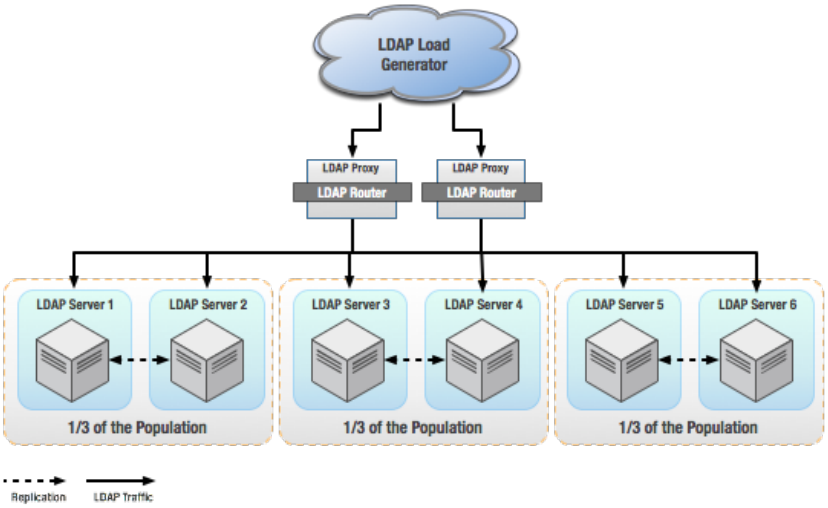

In the case of all other vendors, the SearchMod test used the same flat architecture for Search and Authentication described above with one notable exception; Ping, was also tested with a flat replication across all the LDAP servers using the LDAP proxies fronting and load balancing the traffic. While the use of proxies is not Ping's preferred model for write priority-based services as they add latency, we wanted to test the effects on a similar design when comparing to CA.

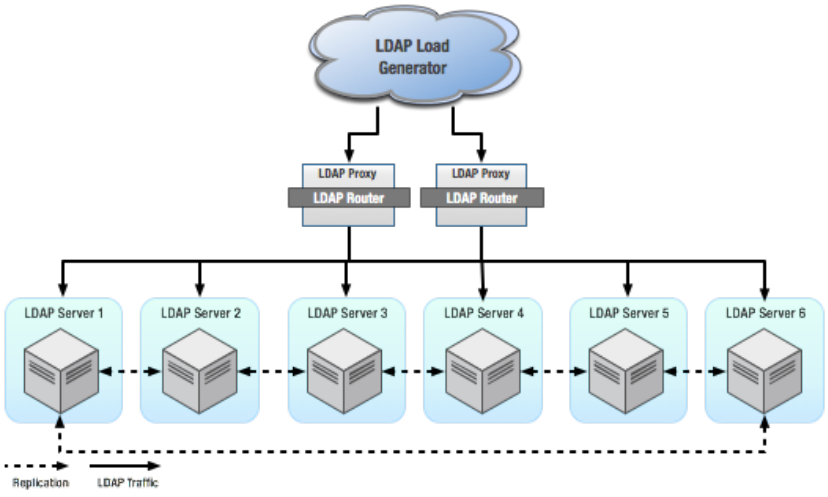

Benchmark Results and Observations

As vendor recommendations and best practices differ based on the write/read priority given to the operations, the results and observations were also split into:

- **Authentication/Search and Compare** - while all read/write traffic was concurrently applied, we focused on the priority traffic that consisted of binding, authenticating, searching and compare operations.
- **Modify/Write** - while all read/write traffic was concurrently applied, we focused on the priority traffic that consisted of binding, searching and modify operations.

Before diving into the split read/write results, the following chart and table provides a summary of the *combined* number of operations per second (given the load mix described in the testing protocol section of this document) that each of the products were able to sustained in support of a large user target population base.

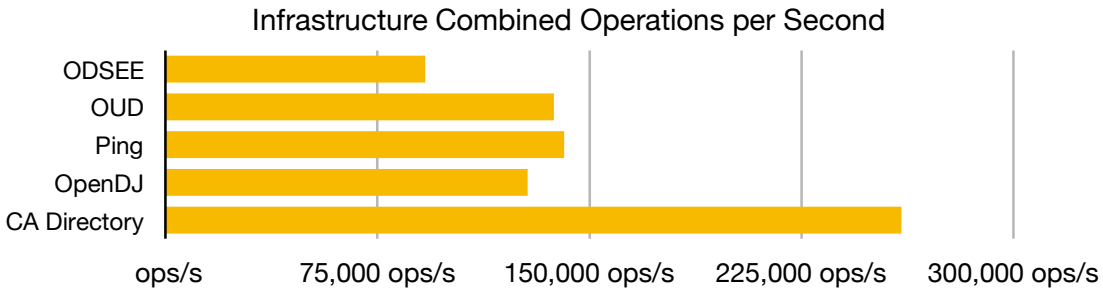

	ODSEE	OUD	Ping	OpenDJ	CA Directory
Total	92,024 ops/s	137,484 ops/s	141,127 ops/s	128,078 ops/s	260,228 ops/s

These results are based on the same, flat test design (no proxy, just the load balancer), performing all the test profiles concurrently while sustaining `0.8ms ~ 0.9ms` latency levels on average (with the exception of ODSEE which was sustaining `1.2ms ~ 1.7ms` latency average for most operations). A noted observation was the average CPU utilization across all the servers during the tests, while CPU utilization remained very high across all products/vendors, the CA LDAP remained within 14% of CPU utilization.

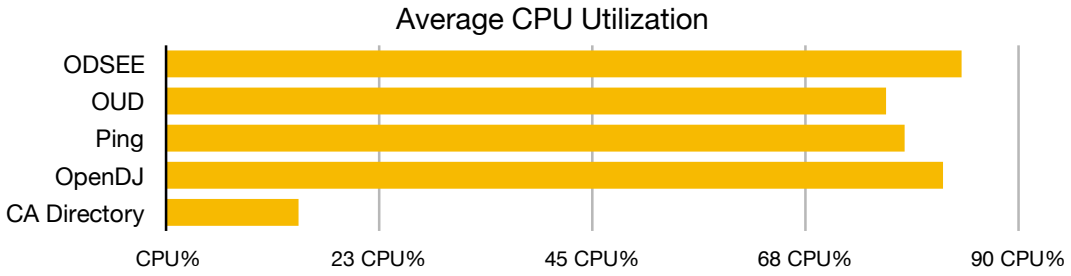

Authentication/Search and Compare Observations

The following charts summarize the results achieved during the Authentication, Search and Compare test across all the products using the same, flat design, direct access to the LDAP servers.

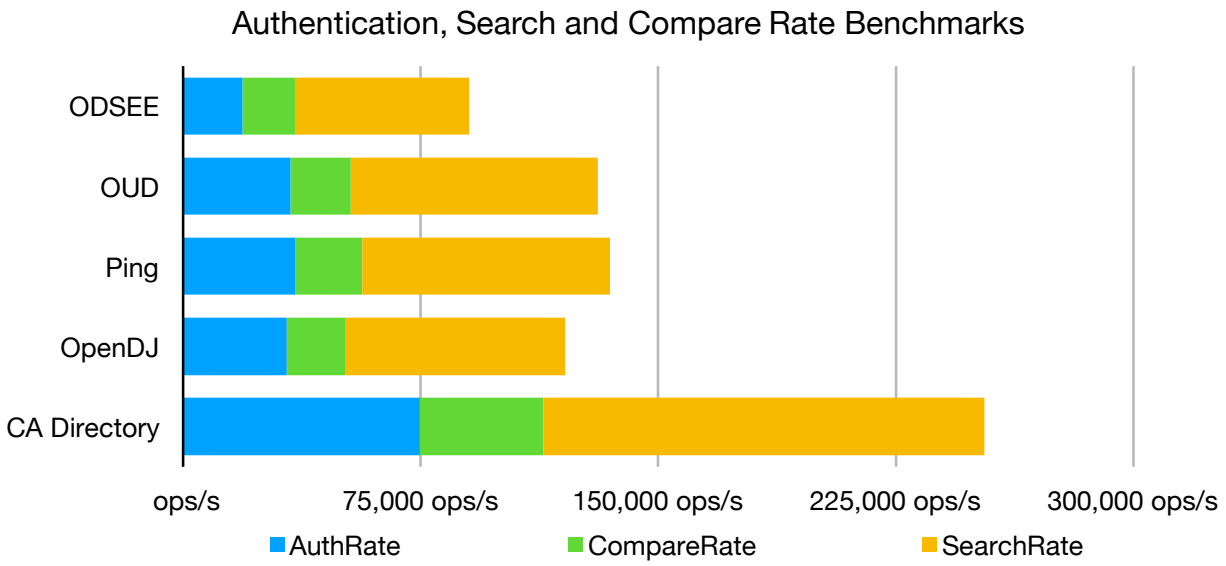

Results of these tests were consistent with historical observations of the different number of operations that LDAP servers are usually able to sustain during mixed operation load. Authentication rates tend to be the lowest as they involve the hashing and compare of the password which takes longer operationally than clear text compare or search rates. As noted by the charts, it is the volume of operations of the CA LDAP that was unique (essentially doubling volume capacity), demonstrating an increased level of performance under heavy load.

Authentication, Search and Compare Benchmarks (with latency)

	ODSEE		OUD		PING		OpenDJ		CA Directory	
	Operations per Second	Average Latency	Operations per Second	Average Latency	Operations per Second	Average Latency	Operations per Second	Average Latency	Operations per Second	Average Latency
AuthRate	18,837 ops/s	2.261ms	33,876 ops/s	1.257ms	35,294 ops/s	1.357ms	32,756 ops/s	1.462 ms	74,637 ops/s	0.637ms
CompareRate	16,766 ops/s	.925ms	19,094 ops/s	0.812ms	21,171 ops/s	0.565ms	18,345 ops/s	0.652ms	39,066 ops/s	0.303ms
SearchRate	54,653 ops/s	1.239ms	77,985 ops/s	0.868ms	78,250 ops/s	0.611ms	69,565 ops/s	0.687ms	139,349 ops/s	0.347ms

Modify/Write Observations

The following charts summarize the results achieved during the Modify/Write test across all the products using the vendor recommended design, direct access to the LDAP servers.

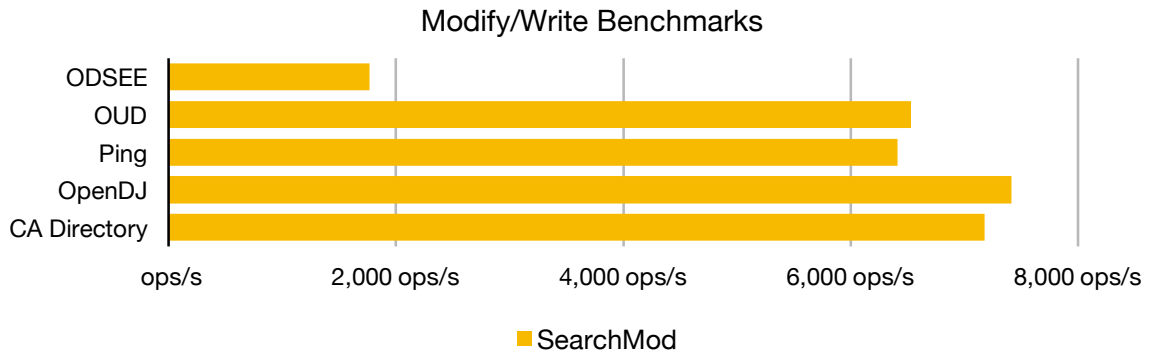

These results reflect the rate to which each of the Directory products could sustain (given the described architecture) a number of modifications per second. However, when examining these write operations it is important to look at the latency of the write transaction to a single server and the overall latency of the write across all replicated servers in the environment. The following table shows the average latency on each user search and each subsequent modify of the selected attribute that each of the SearchMod operation performed.

Modify/Write (Operation Breakdown)

	OUD	Ping	OpenDJ	CA Directory
SearchMod Operation	**6,528 ops/s**	**6,413 ops/s**	**7,412 ops/s**	**7,175 ops/s**
- avg Search Duration (ms)	0.889ms	0.658ms	0.742ms	0.977ms
- avg Modification Duration (ms)	1.953ms	1.803ms	1.410ms	1.248ms
avg Latency per Operation	**2.842ms**	**2.461ms**	**2.152ms**	**2.225ms**

* Each of the SearchMod operations includes the search of the user DN that is to be modified and then the subsequent modification of the attribute. This table shows the breakdown of the latency each of these two activities took in order to perform one (1) SearchMod operation.

* The totals for the Modify/Write were captured for ODSEE but not the breakdown of these operations given the poor performance shown during these tests.

While a conclusion could be drawn that both OpenDJ and the CA Directory were performing comparable Search and Modify rates, CA Directory has an important product feature that sets the CA Directory results apart from the other vendor solutions.

CA Directory averages `2.225ms` to complete the search and write modification including the commit between all servers involved in the replication. This is an important difference and departure from product architecture that other LDAP vendors offer. For instance, the Ping directory service took `2.486ms` latency to complete a similar operation, but only involved the write commitment of that given server before notifying the LDAP client that the operation had concluded. Once that server committed the write request, the (background) replication process attempts to alert all other servers that an update is required in their dataset. Products such as Ping/OpenDJ/OUD display as part of their replication status, the number of write/change commitments “left” for a server to still perform, thus monitoring for replication backlog is an important operational activity.

During these tests, CA Directory was performing 7,175ops/s, where these writes were saved and replicated across *all* the servers that formed part of the Horizontal Partition Configuration before notifying the client that the operation had concluded successfully. In the case of all the other products, these were the results of "one" (1) server committing the modification and then (after the LDAP client had been notified the operation was concluded), replicating that change across the infrastructure in the background. This meant that servers were holding a backlog of as many as 2000~3000 changes at any given moment which were waiting to replicate. It took 6-9 minutes for this queue to drain, meaning that an LDAP search on any of these other servers within the directory bank (potentially as a result of load balancing the searches) would return the original value and not the modified one, presenting the application with an inconsistent view of the data set

Architecturally, services have always had to cope with these small delays by ensuring that the load balancer is always sticking to a server to respond to a subsequent search to ensure that the response is consistent to the previous modification. There are other architectural approaches to solve this data integrity problem in a high traffic environment, but in the case of CA Directory this is addressed inherently as part of the product as data integrity after a write operation is assured by all servers before the LDAP client is notified that the modification has concluded.

Note that in the above results, the test model needed for CA Directory involved the use of the Directory Routers, and architecturally, these are components that add a significant level of latency to an operation.

To compare how CA Directory would fair with another product that also included the use of router/proxies, a second test using Ping was performed. The following chart depicts the results of Ping and CA Directory, both using their respective Router/Proxies to better observe and compare under similar architectural models[1].

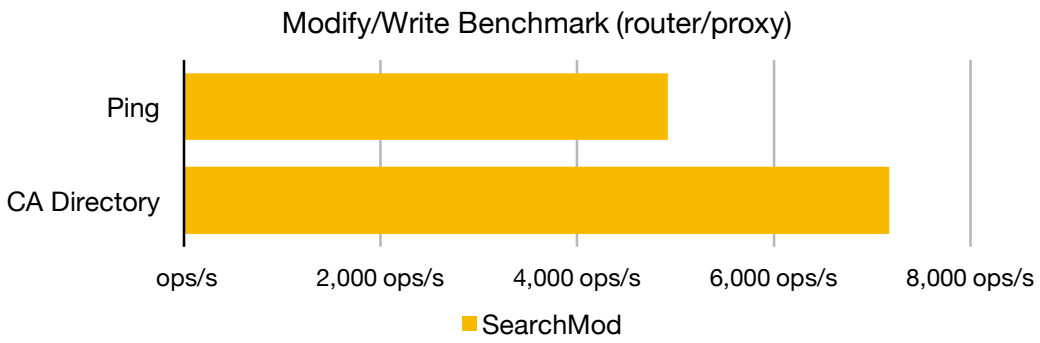

Ping's ability to sustain 4,922ops/s still predicated in the write commit of only one LDAP server (while having the Ping's LDAP proxy handling the load-balancing) while CA Directory's rate of 7,175ops/s includes the write commitment of the servers included within the Horizontal Partition Configuration and the added latency of the CA LDAP Router, which, for the write operations, is a significant increase in the volume of operations when compared to a similarly deployed Ping environment.

Conclusions

This benchmark was conducted to determine if there was a significant difference in performance among the four vendors assessed. It was based on a generic deployment design to simulate a large distributed directory deployment capable of supporting 100 million users.

[1] Ping's proxy does add latency to all operations and it is preferably used when schema partitions or operation re-writes are required as part of the design. This is not necessarily how the vendor would suggest in support of write priority deployments.

The measured volume of operations per second sustained for Oracle, Ping, and ForgeRock were comparable. However, CA Technologies showed 84% higher performance than the other vendors. In addition to gaining higher performance, this translates to a much higher return on an organization's infrastructure investment.