WHITE PAPER | JANUARY 2014

CA Identity Manager One Hundred Million User Test: Results & Analysis

Table of Contents

Section 1:

User test summary and results

In today's environment the types of user communities that an organization must support continue to expand beyond their employee base. Typically, commercial and home-grown Identity Management (IDM) products have been used to support internal employees and third party users who required identity and user access within an organization's IT infrastructure. However today, potentially millions of external users and non-employee identities may need to be registered and managed throughout the entire identity lifecycle. This paper focuses on the requirements and testing of IDM solutions to support a high and growing volume of users throughout the identity lifecycle.

In testing the performance capabilities and scalability support for user populations of up to 100 million user population, the following types of scenarios were considered:

- A government agency allows citizens to self identify and register for access to external facing applications. Potentially millions of citizens may need to register for a specific event, online notification or government-to-citizen account.
- A company selling goods and services over the internet needs to securely capture and manage their customer information for real-time purchasing, enable faster checkout processing and simplify opportunities for repeat business. Usage may be cyclical and highly dependent on specific events, days and times that can cause a spike in user activity.
- Consumer products and retail establishments that need to securely capture and track customer responses to a global promotion. New user registrations will cause an increase in normal volume as the promotion is rolled out to different regions.

The tests and results outlined in this paper are the result of a combined effort between Accenture and CA Technologies in order to determine if the CA Technologies Identity Management solution can support the performance requirements of organizations that must support an increasing volume of both internal and external users. In the end, the IDM solution architected for this testing supported up to 100 million user population and achieved peak performance of 39 transactions per second, demonstrating its ability to adequately scale in order to meet the needs of organizations with an ever increasing volume of employees, partners, suppliers and customers.

1.1 Test

The joint CA Technologies and Accenture test team architected and assembled an environment to simulate the conditions and configuration required for an IDM solution to meet the requirements of a large scale identity proofing and registration system. In this paper, we will describe

- The CA Technologies products leveraged to test performance of the IDM solution
- The hardware and software test environment
- The joint Accenture/CA Technologies test approach for validating the solution's scalability
- Scenarios utilized to mimic real world usage
- Executed test cases and accompanying results

- Lessons learned related to maximizing solution performance

1.2 Results

The solution architected for testing the scalability and performance of CA Identity Manager supported varying user populations of 25, 50 and 100 million users respectively. At a high level the following results were seen across the test cases:

- The solution was able to sustain an average volume of over 30 transactions per second throughout multiple scenarios and test cycles.
- With spikes in the volume and load on the solution an additional burst rate of up to 39 transactions per second was achieved.
- The mix of specific scenarios (including self registration, password reset, self modification, and administrative lock) did not significantly change the results.
- The solution withstood high volumes of users and virtual transactions without major failures or degradation of backend processing with an error ratio of less than 0.38%.

Figure A.

CA Identity Manager 100 Million User Test—Summary of Test Results

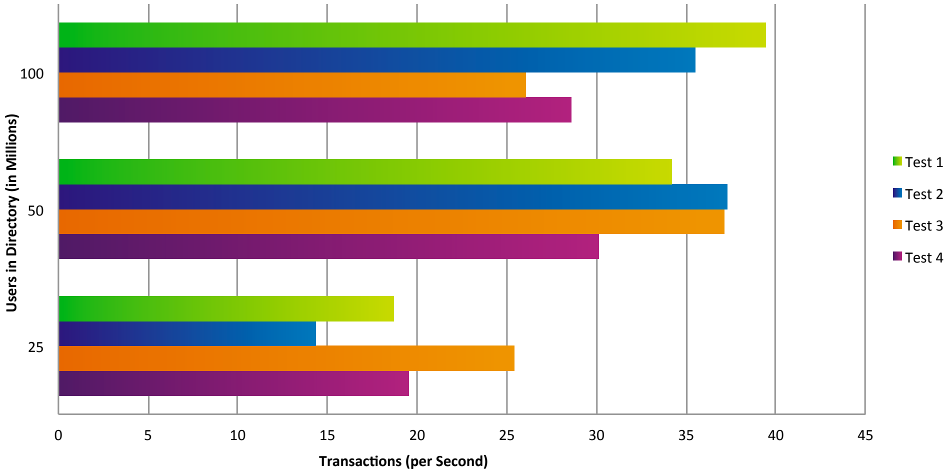

Section 2:

Challenge

With the typical boundaries between internal and external user communities quickly blurring for many organizations, it is critical that the solutions in place to manage and support identities be able to scale in size and complexity to meet the demands of an increasing user base. This has turned a typical IT function of identity management into a value driver and business enabler for the entire organization.

Today many types of organizations are struggling with the ability to handle a large number of user transactions and the interdependency with internal and external facing applications. What used to be the problem for a select group of governmental agencies has evolved into an issue that spans sectors such as Consumer Products, Financial Services, Healthcare, Manufacturing, and almost every other industry.

2.1 Business challenges

Federal and public sector

Local, state and federal government agencies are tasked with keeping up with and trying to meet the demand of a surging consumer base. Identities are not limited to employees or third party contractors but instead extend to every citizen under their jurisdiction. This leads to an expanding user community that Federal and Public Sector organizations are struggling to manage in a cost effective and efficient way. And as the world evolves more and more to instant connectivity, online applications and mobile access to data that demand continues to increase every day.

For example, nearly 99 million individuals filed their US Federal income tax returns electronically during 2010, a 3 percent increase in the IRS e-file rate. Of the 141.5 million returns filed as of this writing in 2011, almost 70 percent were filed electronically.

Each year, more taxpayers chose to e-file their tax returns. In 2009, nearly 95 million taxpayers or 67 percent used e-file. In the past decade, the number of individual tax returns e-filed has increased by 145 percent. The overall number of individual tax returns increased only by 8 percent. IRS e-file is no longer the exception; now it is the norm.
irs.gov/newsroom/article/0,,id=231381,00.html

Commercial sector

Commercial organizations, regardless of industry affiliation, are not immune to the business requirement of knowing and managing users beyond the “four walls” of the organization. From managing customer relationships and personal information to tracking consumer response and spending habits, many significant business needs are driving IT organizations to think about identity management in a whole new way.

For example, e-commerce is here to stay (in Q2011, e-commerce spending exceeded $37B, up 14% from 2010) and companies who sell over the internet need to capture and manage their customer information for real-time purchasing to enable faster checkout times, track user trends and holistically manage their customer’s identity information.
prnewswire.com/news-releases/comscore-reports-375-billion-in-q2-2011-us-retail-e-commerce-spending-up-14-percent-vs-year-ago-127237503.html

2.2 Real-world scenarios

To interpret the findings gathered during the hundred million user testing of the CA Technologies IDM solution it is important to consider real-world usage scenarios. This enables a meaningful comparison of the data points outlined in the test results described in this white paper to show how these results can be extrapolated to meet specific business use cases and requirements. As an example, the table below provides a sample of the potential throughput required in transactions per second for average usage spread across the year and for peak loads during a high volume month.

Figure B.

CA Identity Manager 100 Million User Test—Real-World Rates

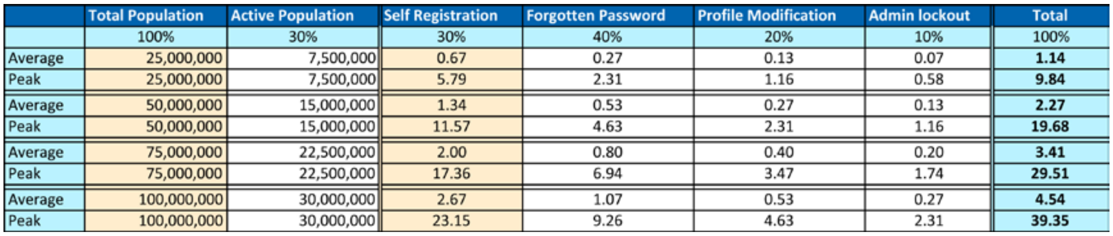

	Total Population	Active Population	Self Registration	Forgotten Password	Profile Modification	Admin lockout	Total
	100%	30%	30%	40%	20%	10%	100%
Average	25,000,000	7,500,000	0.67	0.27	0.13	0.07	**1.14**
Peak	25,000,000	7,500,000	5.79	2.31	1.16	0.58	**9.84**
Average	50,000,000	15,000,000	1.34	0.53	0.27	0.13	**2.27**
Peak	50,000,000	15,000,000	11.57	4.63	2.31	1.16	**19.68**
Average	75,000,000	22,500,000	2.00	0.80	0.40	0.20	**3.41**
Peak	75,000,000	22,500,000	17.36	6.94	3.47	1.74	**29.51**
Average	100,000,000	30,000,000	2.67	1.07	0.53	0.27	**4.54**
Peak	100,000,000	30,000,000	23.15	9.26	4.63	2.31	**39.35**

Note: Throughput for *Average* usage was assumed to be 260 working days throughout the course of an entire year, while *Peak* usage was condensed to a single 30 day window with the same population. Calculations assume 12 hours of activity per day across the total number of days for each row. For both *Average* and *Peak*, the transactions per second were calculated against the *Total Population* for *Self Registration* (highlighted in yellow) and the *Active Population* for *Forgotten Password*, *Profile Modification* and *Administrative Lockout*.

Section 3:

CA Technologies solution

CA Technologies provides an Identity Management solution to meet the business challenges created by a 100 million user population. The solution comprises CA Identity Manager and CA Directory. When combined, these two products can provide a scalable Identity Management solution capable of processing high volume transactions without significant performance degradation.

3.1 CA Identity Manager

Identity Management solutions are a key component for managing increasing numbers of identities in any organization. These solutions can automate the business processes so that an employee receives the required accesses to information systems without compromising security and increases end-user productivity and satisfaction.

CA Identity Manager provides an integrated identity administration solution, serving as the foundation for user provisioning, self-service requests, identity governance, and other key processes, which enables your organization to:

- Automate the processes of on-boarding, modifying, and off-boarding users and their associated access.
- Enable end users to initiate provisioning actions, ongoing password management, as well as delegate administration and related processes.
- Leverage role and policy analysis, and certification processes with CA Role & Compliance Manager for more accurate and efficient provisioning and governance decisions.
- Quickly adapt policies and applications when the business changes by utilizing powerful wizards and graphical tools—i.e. PolicyXpress—eliminating the need for custom coding.

Benefits

By establishing a strong identity management infrastructure, CA Identity Manager enables your organization to:

- Detect non-compliant provisioning activities before they happen.
- Eliminate guesswork when providing access to users based on their roles.
- Enable user-based access requests and approvals without involving the help desk.
- Provide employees, contractors, and partners with access to the applications they need on the day they start, and remove it immediately when terminated.
- Adjust the identity management infrastructure easily whenever business processes change.

Features

CA Identity Manager contains a wide range of features for managing identities and access rights and meeting identity compliance requirements.

- **User provisioning & deprovisioning.** Automates account provisioning, removal, and approval processes throughout the user's lifecycle: from on-boarding to fine-grained entitlement management to role changes and ultimately to off-boarding. Customizable workflows support the unique way each organization approves, alerts, and schedules these activities.
- **User self-service.** Enables users to manage attributes of their own identities, easing the burden on IT or the help desk and empowering the business to make identity-related decisions. Users can reset passwords, request access to resources according to organization defined approval workflows, and manage both profile and security information.
- **Centralized control of users, roles, and policies.** Aggregates and synchronizes information across multiple identity silos to manage user attributes, identity tasks, and roles in accordance with central security controls. CA Identity Manager provides an out-of-the-box interface which can be customized to the look, feel, and function needed by each organization or user.

- **Customization without custom code.** Out-of-the-box integration with leading target systems—combined with powerful features such as Policy Xpress, Connector Xpress, a graphical workflow editor, and ConfigXpress—lets you customize your identity management infrastructure to your unique organizational requirements without writing, managing, and supporting custom code.
- **Deep cross-IAM integration.** CA Identity Manager delivers out-of-the-box integration with other IAM solutions from CA Technologies, including CA Role & Compliance Manager, CA User Activity Reporting Module, CA DLP, and CA SiteMinder®. These integrations further extend the value of CA Identity Manager and enhance existing CA IAM implementations.
- **Securing on-premises and cloud applications.** Provides a single point of management for identities across enterprise applications, whether residing on-premise in mainframe or distributed systems or hosted in the cloud. CA Identity Manager provides out-of-the-box connectors for leading applications such as SAP, CA ACF2™, Active Directory, and Salesforce.com

For more information on CA Technologies solution for Identity Management and Governance, visit ca.com/us/secure-identity.aspx

3.2 CA Directory

Enterprise directory solutions must have exemplary performance, scalability and reliability. CA Directory meets these requirements with its ability to mesh any number of servers (including external LDAP resources) into a backbone infrastructure capable of seamless and transparent distributed operations, guaranteed consistent replication, and automatic recovery.

CA Directory:

- Provides scalability without high hardware costs.
- Meets the needs of new, dynamic business applications.
- Improves operational efficiency by consolidating islands of data into a single information backbone.
- Provides a highly responsive and a highly available experience for online application users.

DXgrid

Integral to CA Directory is DXgrid, which is a revolution in directory systems. Compared to the typical LDAP database approach, DXgrid allows for unprecedented levels of scalability, reliability and performance. DXgrid achieves its performance by using a memory-mapped store and, as seen in testing, this can cut the time to bulk load records from days to hours, while at the same time drastically minimizing hardware requirements. The memory-mapped file loaded on each DSA in the CA Directory deployment during start up minimizes disk I/O. The technology includes both shortest path routing and parallel search capability between cooperating servers. Reliability is also enhanced by using write-through rather than write-behind technology, and a safe load-sharing and failover capacity.

Section 4:

Test Environment

The test environment for the 100 million user test was stood up on the Amazon Elastic Compute (EC2) Cloud. Amazon Elastic load balancing was used for sharing the transaction load between four CA Identity Manager servers.

4.1 Amazon Elastic Compute Cloud (EC2)

Amazon EC2 is a web service that provides resizable compute capacity in the cloud. It is designed to make web-scale computing easier for developers. Amazon EC2 allows customers to rent virtual machines by the hour to run their own applications.

Amazon EC2:

- **Elastic**—Capacity can be increased or decreased within minutes and multiple servers can be commissioned simultaneously
- **Flexible**—Amazon EC2 offers multiple instance types, operating systems and software packages. It offers different options for memory, CPU and storage
- **Reliable**—Amazon EC2 offers a highly reliable environment where replacement servers can be deployed rapidly and predictably commissioned. Amazon EC2 Service Level Agreement commitment is 99.95% availability for each region
- **Secure**—Amazon EC2 provides numerous ways of securing resources including web services interfaces for configuring firewall rules as well as multiple layers of protection and controls around the hypervisor and physical protections in the data center.

For more information on Amazon EC2 visit amazon.com/ec2/

4.2 Architecture

For the corporate store architecture, CA Directory was deployed, consisting of 6 directory servers, configured into 3 pairs. The directory was configured into a horizontal partition. The total user population was deployed across a 6 node deployment. Each node was responsible for about 1/6 of the user population.

Figure C.

CA Identity Manager 100 Million User Test— Test Environment

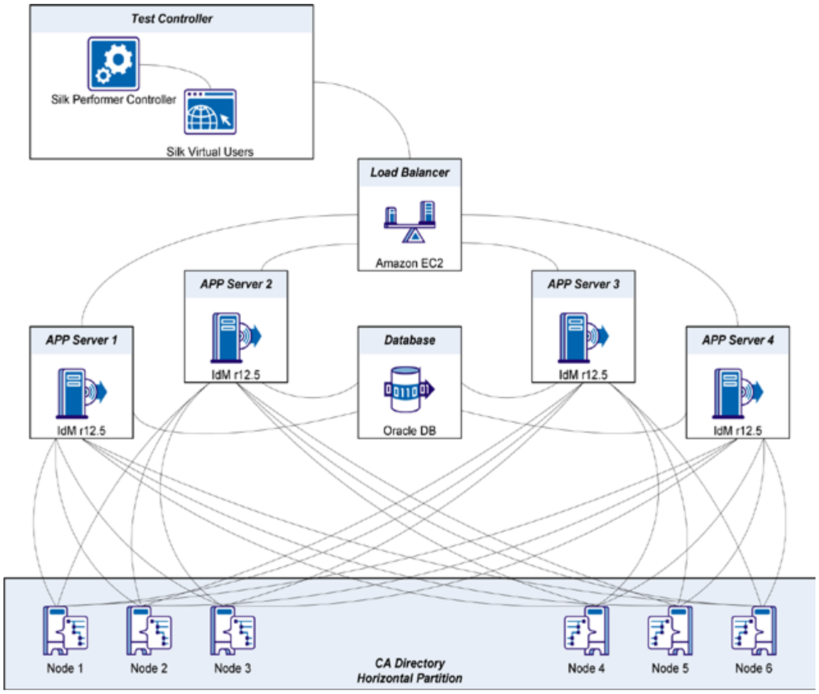

CA Identity Manager was deployed on JBOSS on 64 bit CentOS servers. Router DSAs were configured on the JBOSS servers for CA Identity Manager to communicate to the CA Directory infrastructure. The back end database was on Oracle 11G on Oracle Enterprise Linux. The Silk Test controller was installed and configured on Windows 2008 Server.

4.2.1 Application specs

Solution Component	Hardware Used	Number of Systems
Application Tier	▪ CA Identity Manager 12.5.7.0.580 ▪ JBOSS 5.1.0 Slim (Clustered)	4
Directory Tier	▪ CA Directory r12.0 build 4574	6
Database Tier	▪ Oracle 11g (11.2.0.1.0)	1

4.2.2 Hardware specs

Solution Component	Hardware Used	Number of Systems
Application Tier	▪ RHES 5.4 (Cent OS) ▪ 4 core (2.67 GHZ) ▪ 34.2 GB RAM ▪ 100 GB HD	4
Directory Tier	▪ RHES 5.4 (Cent OS) ▪ 8 core (2.67 GHZ) ▪ 68.4 GB RAM ▪ 100 GB HD	6
Database Tier	▪ Oracle Enterprise Linux ▪ 8 core (2.67 GHZ) ▪ 68.4 GB RAM ▪ 100 GB HD	1

Section 5:
Tests

The tests for this project were modeled after the characteristics of a typical B2C environment. The B2C profile will have more user centric self-help activities like registration and forgotten password.

The test scenarios were intended to represent a large B2C public web site that currently has 100,000,000 users. We assumed that users were being provisioned to a single large user directory.

Identity management tasks were executed using a pool of virtual test users where the virtual test user executed different tasks that are representative of and in proportion to a B2C IDM deployment.

It is important to stress that in an IDM deployment, not all users utilize the identity management system each day and as a result, a model was developed for this test identifying the number of each of the identity management operations expected to take place every 90 days and through that number, a ratio of tasks and a target of concurrent operations was established.

5.1 Use cases

The use cases tested in the 100 million user test were modeled after a B2C scenario, where most of the self-service tasks are used predominantly.

- **Self registration.** This task enables a user to register to a web site after answering a few questions. On submission of the task , the user receives a confirmation and an identity for the user is created in the user store.

- **Forgotten password reset.** This task allows a user to reset their password if they forget their password. CA Identity Manager verifies the identity of the user by presenting verification questions to the user and at the end of the verification allows the user to reset his password.
- **Profile modification.** Allows a user to modify the information about him in the directory. The user has to authenticate and a profile page will be presented where information can be updated.
- **Administrator account lock.** Allows an administrator to enable or disable a user's account.

5.2 Test approach/plan

The goal of the 100 million user testing was to determine the performance of the CA Identity Manager for a large user population. A test execution plan was developed to stress the system with a mix of self-service tasks and varying loads.

- Three test sets were executed for this project based on the number of entries in the user directory. The user directory was populated with 25 million, 50 million and 100 million users for the test sets.
- Each test set consisted of four different tasks: Self Registration, Forgotten Password Reset, User Self Modification, and Administrator Account Lock. As the tests progressed, the tasks were mixed together to replicate a real world load.
- All the tests were run for 20 minutes.

The task mix for the different tests is shown in the table below.

Test	Task mix for the different test cycle (In percentage of users)
Test 1	Self Registration—100% of the users
Test 2	Self Registration—80%, Forgotten password—20%
Test 3	Self Registration—70%, Forgotten Password—20% User Self Modification—10%
Test 4	Self Registration—60%, Forgotten Password—20% User Self Modification—10%, Administrator Account lock—10%

5.2.1 Test 1

For test cycle #1, all the virtual users were assigned to perform self-registration activities. The tests were run for twenty minutes with an average of 31 transactions per second and a maximum throughput of 39 transactions per second.

Figure D.

CA Identity Manager 100 Million User Test Results—Transactions per Second with User Self-Registration

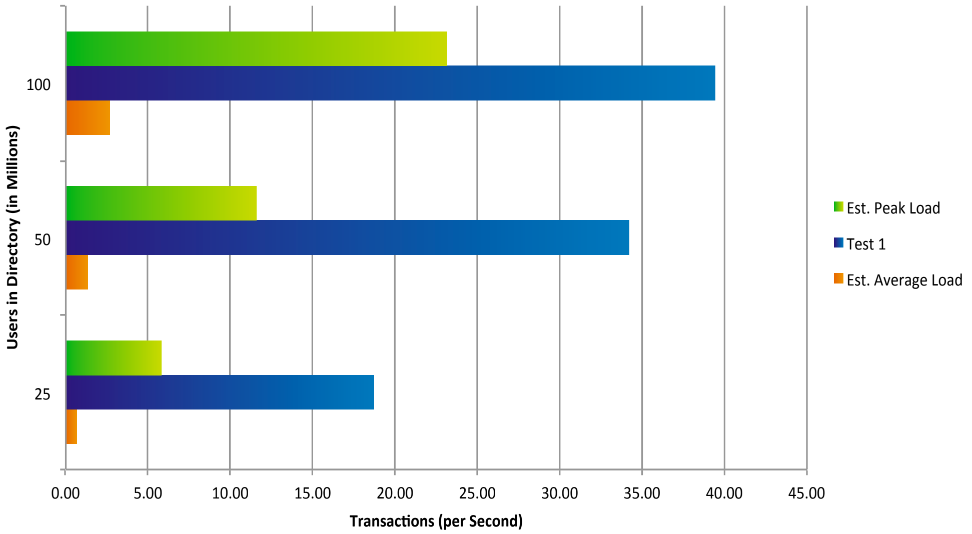

5.2.2 Test 2

For test cycle #2, the virtual users were split between self-registrations (80%) and forgotten password (20%) activities. The tests were run for twenty minutes with an average of 29 transactions per second and a maximum throughput of 37 transactions per second.

Figure E.

CA Identity Manager 100 Million User Test Results—Transactions per Second with User Self-Registration and Forgotten Password

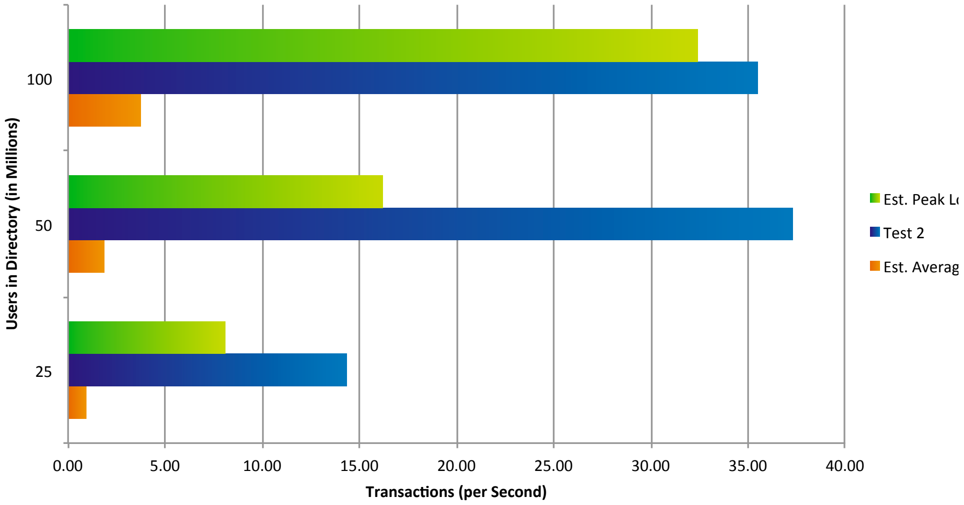

5.2.3 Test 3

For test cycle #3, the virtual users were split between self-registrations (70%), forgotten password (20%), and user self-modification (10%) activities. The tests were run for twenty minutes with an average of 30 transactions per second and a maximum throughput of 37 transactions per second.

Figure F.

CA Identity Manager 100 Million User Test Results—Transactions per Second for User Self-Registration, Forgotten Passwords, and User Self-Modification

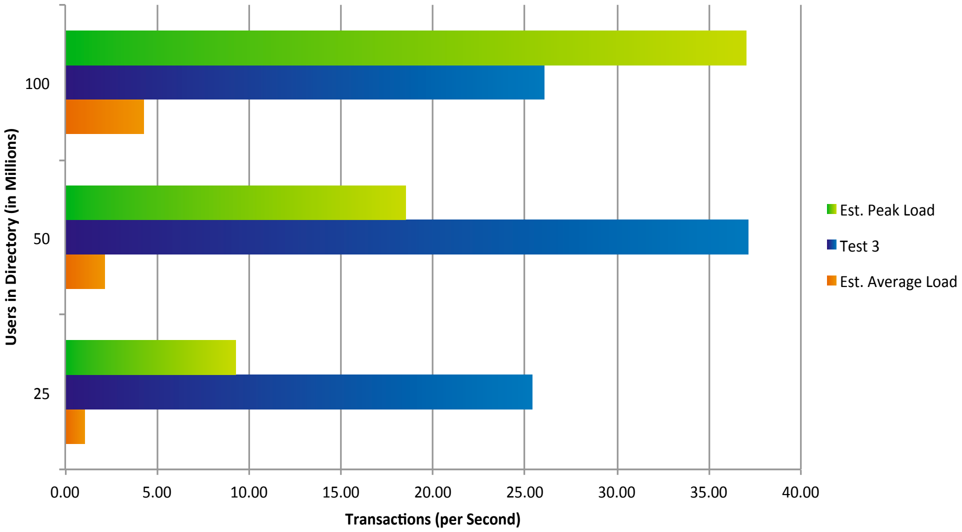

5.2.4 Test 4

For test cycle #3, the virtual users were split between self-registrations (60%), forgotten password (20%), user self-modification (10%), and administrator account lock (10%) activities. The tests were run for twenty minutes with an average of 26 transactions per second and a maximum throughput of 30 transactions per second.

Figure G.

CA Identity Manager 100 Million User Test Results—Transactions per Second for User Self-Registration, Forgotten Passwords, User Self-Modification and Admin Account Lock

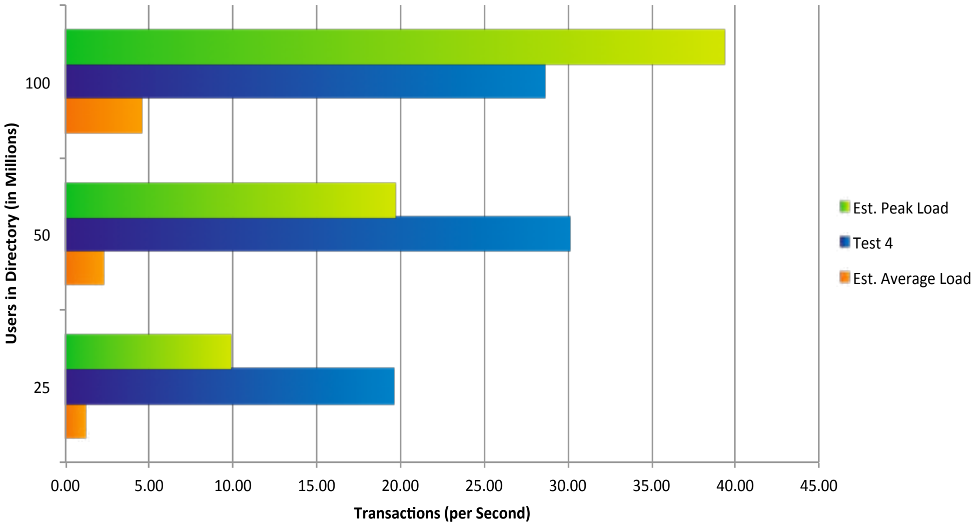

5.3 Tuning considerations

Overall performance largely was not impacted by the number of users housed in the underlying data store. Instead tuning made at the Application Server, Directory and Database level helped to increase performance around the number of transactions able to be processed by the IDM solution.

Application Server

- By adding additional server clustering a higher transactional throughput rate can be achieved by scaling the deployed solution horizontally across the environment.
- In order to increase throughput, the Java heap size was adjusted to increase performance on the Identity Manager product's ability to process transactions.
- To reduce overhead at the application layer logging was kept to a minimum.
- Increased settings at the application tier to properly throttle the connection ties to the database in order to allow additional throughput between the application server and database tiers.

Directory

- Scaling the directory server by leveraging horizontal partitioning at the directory layer added redundancy at the directory level for the architecture designed for this testing effort.

Database

- Oracle 11g was used for the database server in order to improve performance and provide a higher transactional read/write throughput at the database layer.
- Additional fine-tuning and configuration adjustments on the database connections as well as management of table spaces for task persistence yielded higher performance results across test cycles.
- Clustering of the database can provide a better overall performance and remove the single bottleneck found in the architecture used during testing.

Section 6:

Conclusion

Varied iterations of standard identity management use cases were utilized to simulate test cycles and mimic scenarios that would address a range of business requirements. These tests demonstrated that CA Identity Manager is capable of performing and scaling to support up to a 100 million user population. And when the data found during the joint Accenture/CA Technologies testing (average throughput of 30 transactions per second and peak loads of 39 transactions per second) are compared to the example real-world scenarios described in the business case section of this whitepaper (average throughput range of 1–4 transactions per second and peak loads of 9–39 transactions per second), the results show that the IDM solution architected for these tests could meet or exceed these potential requirements for throughput ratios in a high volume environment.

For more information about Accenture, visit accenture.com/security

For more information about CA Identity Manager, go to ca.com/identityminder ca technologies

Connect with CA Technologies at ca.com

Agility Made Possible: The CA Technologies Advantage

CA Technologies (NASDAQ: CA) provides IT management solutions that help customers manage and secure complex IT environments to support agile business services. Organizations leverage CA Technologies software and SaaS solutions to accelerate innovation, transform infrastructure and secure data and identities, from the data center to the cloud. CA Technologies is committed to ensuring our customers achieve their desired outcomes and expected business value through the use of our technology. To learn more about our customer success programs, visit **ca.com/customer-success**. For more information about CA Technologies go to **ca.com**.