# Access Manager

## Performance and Sizing Guidelines

**March 2021**

## Legal Notice

For information about legal notices, trademarks, disclaimers, warranties, export and other use restrictions, U.S. Government rights, patent policy, and FIPS compliance, see https://www.microfocus.com/about/legal/.

**© Copyright 2021 Micro Focus or one of its affiliates.**

# Contents

# About this Book and the Library

This guide provides the performance and sizing recommendations for Access Manager. This information helps you in deploying the correct configuration in your environment. The test results are simulated in a lab environment.

On similar hardware, you may have different results. The test result may vary based on the applications used, type of data, user store, and a number of other dependent components operating in the environment. It is recommended to first verify the performance in your environment before deploying the product in a high-scale environment.

For information about the test strategy, hardware, and software used in the tests, Section A.1, “Test Strategy,” on page 27.

## Other Information in the Library

You can access other information resources in the library at the following locations:

- Access Manager Product Documentation
- Access Manager Developer Resources

---

**NOTE:** Contact namsdk@microfocus.com for any query related to Access Manager SDK.

---

# 1 Access Manager Performance



## 1.1 Access Gateway Appliance

The following performance numbers are recorded in minutes to show how the system performs:

| Test Scenario | Results |
|---|---|
| HTTPS Public (a user accessing single page in a session) | 1700K requests per minute with a throughput of 2000 Megabits per minute |
| HTTPS Public (a user accessing 10 pages in a session) | 1400K requests per minute with a throughput of 5000 Megabits per minute |
| HTTPS Authentications using secure name/password - form | 42K logins per minute |
| HTTPS Authorizations | 30K authorized pages per minute |
| HTTPS Authorization with 10 page requests | 150K authorizations per minute |

The following performance numbers are recorded in seconds to show how the system performs:

| Test Scenario | Results |
|---|---|
| Concurrent Sessions in a 4-node Access Gateway cluster | 240K sessions in cluster (approximately 60K sessions per server) |
| Concurrent Sessions in a 4-node Identity Server cluster | 240K sessions in cluster (approximately 60K sessions per server) |
| HTTP Public | 35K requests per second |
| HTTPS Public | 28K requests per second |
| HTTPS Authentications using Name/Password – Basic | 700 logins per second |
| HTTPS Authentications using Secure Name/Password – Basic | 700 logins per second |
| HTTPS Authentications using Name/Password – Form | 700 logins per second |
| HTTPS Authentications using Secure Name/Password – Form | 700 logins per second |

| Test Scenario | Results |
|---|---|
| HTTPS Login with Roles/Access Gateway Authorization | 500 logins per second |
| HTTPS Login with Identity Injection | 425 logins per second |
| HTTPS Login with Form Fill | 350 logins per second |
| HTTPS Authorizations with 10 page requests | 2500 authorized pages per second |

## 1.2 Access Gateway Service on SLES 12

These test results are for Access Gateway 4.2 on SLES 12. With Access Gateway 4.4, a significant improvement in the performance of Access Gateway for public request has been made. Refer to Access Gateway 4.4 Performance.

The following performance numbers are recorded per minute to show how the system performs:

| Test Scenario | Results |
|---|---|
| HTTPS Public (a user accessing a single page in a session) | 2600K requests per minute with a throughput of 2700 Megabits per minute |
| HTTPS Public (a user accessing 10 pages in a session) | 1600K requests per minute with a throughput of 6200 Megabits per minute |
| HTTPS Authentications using secure name/ password - form | 39K logins per minute |
| HTTPS Authorizations | 30K authorized pages per minute |
| HTTPS Authorization with 10 page requests | 150K authorizations per minute |

The following performance numbers are recorded in seconds to show how the system performs:

| Test Scenario | Results |
|---|---|
| Concurrent Sessions in a 4-node Access Gateway cluster | 260K sessions in the cluster (approximately 65K sessions per server) |
| Concurrent Sessions in a 4-node Identity Server cluster | 280K sessions in the cluster (approximately 70K sessions per server) |
| HTTP Public | 37K requests per second |
| HTTPS Public | 43K requests per second |
| HTTPS Authentications using Name/Password – Basic | 700 logins per second |
| HTTPS Authentications using Secure Name/Password – Basic | 650 logins per second |
| HTTPS Authentications using Name/Password – Form | 650 logins per second |
| HTTPS Authentications using Secure Name/Password – Form | 650 logins per second |
| HTTPS Login with Roles/Access Gateway Authorization | 500 logins per second |

| Test Scenario | Results |
|---|---|
| HTTPS Login with Identity Injection | 400 logins per second |
| HTTPS Login with Form Fill | 450 logins per second |
| HTTPS Authorizations with 10 page request | 2500 authorized pages per second |

## 1.3 Access Manager Appliance

The following performance numbers are recorded per minute to show how the system performs:

| Test Scenario | Results |
|---|---|
| HTTPS Public (a user accessing a single page in a session) | 2808K requests per minute with a throughput of 3000 Megabits per minute |
| HTTPS Public (a user accessing 10 pages in a session) | 1800K requests per minute with a throughput of 6600 Megabits per minute |
| HTTPS Authentications using secure name/password - form | 33K logins per minute |
| HTTPS Authorizations | 24K authorized pages per minute |
| HTTPS Authorization with 10 page requests | 168K authorizations per minute |

The following performance numbers are recorded per second to show how the system performs:

| Test Scenario | Result |
|---|---|
| Concurrent Sessions in a 4-node Access Gateway cluster | 560K sessions in the cluster (approximately 140K sessions per server) |
| Concurrent Sessions in a 4-node Identity Server cluster | 720K sessions in the cluster (approximately 180K sessions per server) |
| HTTP Public | 48K requests per second |
| HTTPS Public | 47K requests per second |
| HTTPS Authentications using Name/Password – Basic | 650 logins per second |
| HTTPS Authentications using Secure Name/Password – Basic | 660 logins per second |
| HTTPS Authentications using Name/Password – Form | 550 logins per second |
| HTTPS Authentications using Secure Name/Password – Form | 560 logins per second |
| HTTPS Login with Roles/Access Gateway Authorization | 400 logins per second |
| HTTPS Login with Identity Injection | 300 logins per second |
| HTTPS Login with Form Fill | 290 logins per second |
| HTTPS Authorizations with 10 page request | 2800 authorized pages per second |

## 1.4 Identity Server as an OAuth 2.0 Identity Provider

The following table lists different OAuth requests on a single node Identity Server and performance for each request:

| Test | Scenario | Access Manager Performance |
|---|---|---|
| Client credentials flow without a refresh token | Users request for an access token in the client credentials flow without a refresh token. | 820 tokens per second |
| Client credentials flow with a refresh token | Users request for an access token in the client credentials flow along with a refresh token. | 800 tokens per second |
| Resource owners flow without refresh tokens | Users request for an access token in the resource owners flow without requesting for a refresh tokens. | 600 tokens per second |
| Resource owners flow with refresh tokens | Users request for an access token in the resource owners flow with refresh tokens. | 200 tokens per second |
| Authorization code flow without refresh tokens | Authenticate and request for an authorization code and using the authorization code request for an access token without requesting for refresh tokens. | 120 tokens per second |
| Authorization code flow with refresh tokens | Authenticate and request for an authorization code and using the authorization code request for an access token with refresh tokens. | 110 tokens per second |
| Implicit flow – access tokens | Request for an access token in the implicit flow. | 140 tokens per second |
| Implicit flow – ID tokens | Request for the ID token in implicit flow. | 140 token per second |
| Implicit flow – Access token + ID tokens | Request for an access token and an ID token in the implicit flow. | 130 tokens per second |
| Token validation | Validate an access token against the tokeninfo endpoint. | 540 validations per second |
| Token refresh | Getting an access token by submitting the refresh token. | 460 token refreshes per second |
| User Attributes | Fetching the user attributes against the userinfo endpoint | 540 requests per second |

For information about the test environment, see Test Environment: Identity Server as an OAuth 2.0 Identity Provider.

**NOTE:** To improve the performance of OAuth requests, scale Access Manager components horizontally by adding additional components to the cluster.

## 1.5 Advanced Session Assurance

This section explains the performance test results and performance impact of enabling Advanced Session Assurance on authentication and sessions in Identity Server and Access Gateway.

- **Accessing Login Pages for the First Time:** A delay occurs when a user accesses the login page for the first time after enabling Advanced Session Assurance. This delay is due to the downloading of initial login pages and associated Java Scripts to the browser. The subsequent requests will not be delayed as Java Scripts are cached in the browser during the first attempt.
- **Advanced Session Assurance Parameters:** An additional delay may occur in the login performance if the following parameters are enabled for session assurance.
  - HTML5 Capabilities
  - System Fonts
  - Webgl Metadata

This delay is due to the client side browser processing for the additional parameters. These parameters do not impact the server side processing.

- Impact of Enabling Advanced Session Assurance on Identity Server Performance
- Impact of Enabling Advanced Session Assurance on the Access Gateway Performance

---

**NOTE:** For information about the test environment, see Test Environment: Advanced Session Assurance.

---

### 1.5.1 Impact of Enabling Advanced Session Assurance on Identity Server Performance

**Use case:** A user logs in to Identify Server ([https://<idp_url>/nidp/app/) continuously with the Secure Name Password form contract with Advanced Session Assurance configured at Identity Server.

| | With Session Assurance | Without Session Assurance | Performance Impact |
|---|---|---|---|
| Logins Per Second | 230 logins per second | 250 logins per second | 8% |
| Number of Sessions | ~200 K sessions | ~200 K Sessions | 0 |

### 1.5.2 Impact of Enabling Advanced Session Assurance on the Access Gateway Performance

**Use case:** A user accesses the protected resource with the secure name password form contract when Advance Session Assurance is enabled for both Identity Server and Access Gateway.

| | With Session Assurance | Without Session Assurance | Performance Impact |
|---|---|---|---|
| Access Gateway requests Per Second | 130 requests per second | 160 requests per second | 18% |

| | With Session Assurance | Without Session Assurance | Performance Impact |
|---|---|---|---|
| Number of Sessions | ~200 K sessions | ~200 K Sessions | 0 |

## 1.6 Components Scalability

The goal of the scalability tests is to validate the architecture and show the size of clusters/ components used.

| Component | Number of Devices |
|---|---|
| Identity Servers | 12 |
| Access Gateway Appliance | 18 |
| Linux Access Gateways | 8 |
| LDAP Servers | 8 |
| Web Servers | 101 |
| Policies/Roles | 101 |
| Accelerators | 51 |
| Concurrent Users on Access Manager | 40000 sessions per Access Gateway |

# 2 Sizing Guidelines

The following recommendations are based on the test results:

- If your environment demands a large number of users active throughout the day, scaling the memory to hold these user sessions is recommended.

  For example, all users keep their portal or mailbox open throughout the day. However, if the environment demands more users activities such as logins per second / requests per second, scaling the CPU is recommended for faster processing. For example, trading systems where a large number of users log in at the same time and leave the session quickly.
- A total number of users in the LDAP user store is not significant for determining hardware for Access Manager components.
- When usage is high on accessing web servers and applications, more Access Gateways are required.
- When usage is high on users and authentication, more Identity Servers are required.
- Two nodes in a cluster are given as the minimum recommended configuration for fault tolerance.
- The setup needs to be evaluated in a real-world usage of the use case.

**In this Chapter**



## 2.1 Recommendation based on Logins per Second

| Logins per Second | Number of Nodes | Server Configuration of Each Node |
|---|---|---|
| Less than 200 | 2 Identity Server<br>2 Access Gateway | 4 X CPU, 16 GB Memory |
| 200 - 500 | 4 Identity Server<br>4 Access Gateway | 4 X CPU, 16 GB Memory |
| 500 - 650 | 6 Identity Server<br>6 Access Gateway | 4 X CPU, 16 GB Memory |

## 2.2 Recommendation based on Active Sessions

| Active Sessions | Number of Nodes | Server Configuration of Each Node |
|---|---|---|
| Less than 200,000 | 2 Identity Server and 2 Access Gateway | 2 X CPU, 16 GB Memory |
| 200,000 – 300,000 | 2 Identity Server and 2 Access Gateway | 4 X CPU, 32 GB Memory |
| 300,000 – 400,000 | 4 Identity Server and 4 Access Gateway | 4 X CPU, 32 GB Memory |

## 2.3 Recommendation based on Access Gateway Hits per Second

| Hits Per Second | Number of Nodes | Server Configuration of Each Node |
|---|---|---|
| Less than 10,000 | 2 Access Gateways | 2 X CPU, 16 GB Memory |
| 10,000 – 20,000 | 2 Access Gateways | 4 X CPU, 16 GB Memory |
| 20, 000 – 40,000 | 2 Access Gateways | 8 X CPU, 16 GB Memory |

## 2.4 Horizontal and Vertical Scaling

These tests include the following Access Manager operations:

- Logins (See Login Performance)
- Active Sessions (See Scalability of Active Users Sessions)
- Hits (See Access Gateway Hits Scalability)
- Throughput (See Access Gateway Throughput Scalability)

**NOTE:** For more information, see Test Environment: Vertical and Horizontal Scalability.

### 2.4.1 Login Performance

This scope of this test is measuring the login performance when a user accesses the resource protected with the Secure Name Password Form contract.


- "Login Performance with CPU Scaling" on page 15
- "Login Performance with Memory Scaling" on page 16
- "Login Performance with Number of Nodes in the Cluster" on page 17

## Login Performance with CPU Scaling

In this test, memory is kept constant at 32 GB and Tomcat is assigned with 16 GB in Identity Server and Access Gateway. CPUs are increased in the following order 1, 2, 4, 8, and 16 and performance is measured at each CPU level.

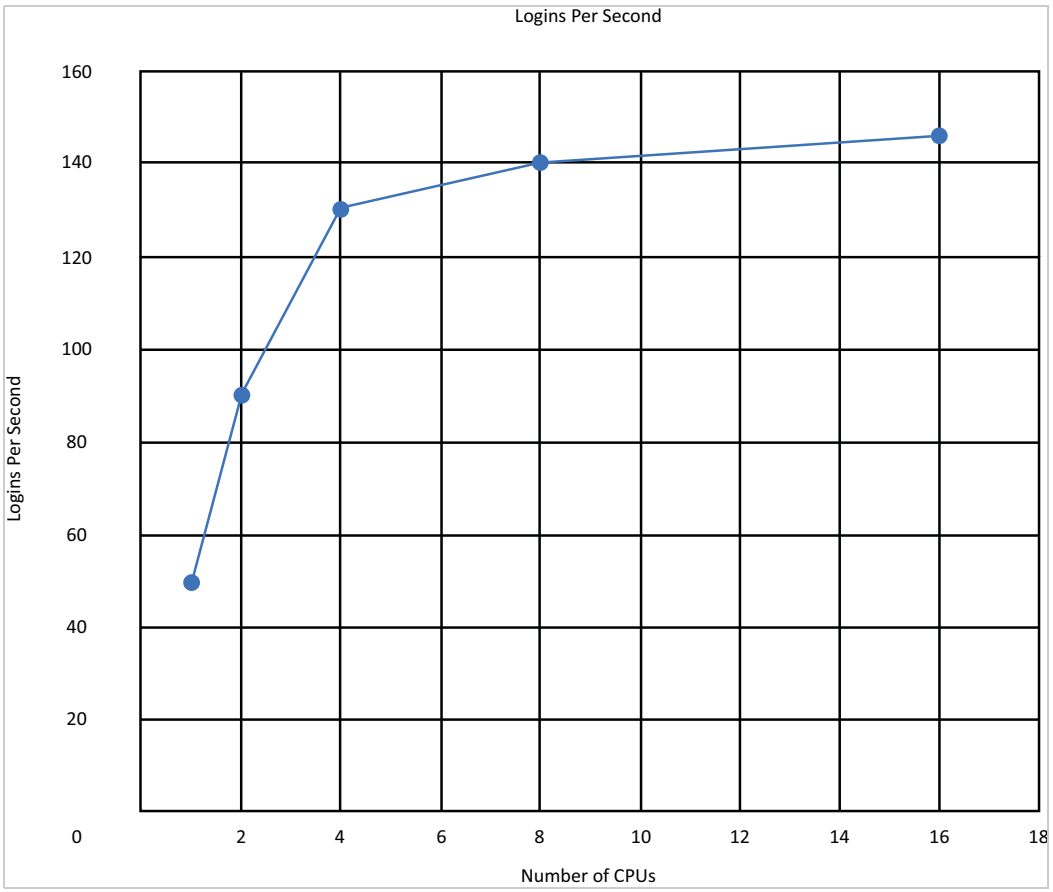

## Login Performance with Memory Scaling

In this test, the number of CPU is kept constant at 16 for Identity Server and Access Gateway. Memory is increased in the order 8 GB, 16 GB, 32 GB, and 64 GB. Also, Tomcat is assigned with 70% of the available memory. Performance is measured at each memory level.

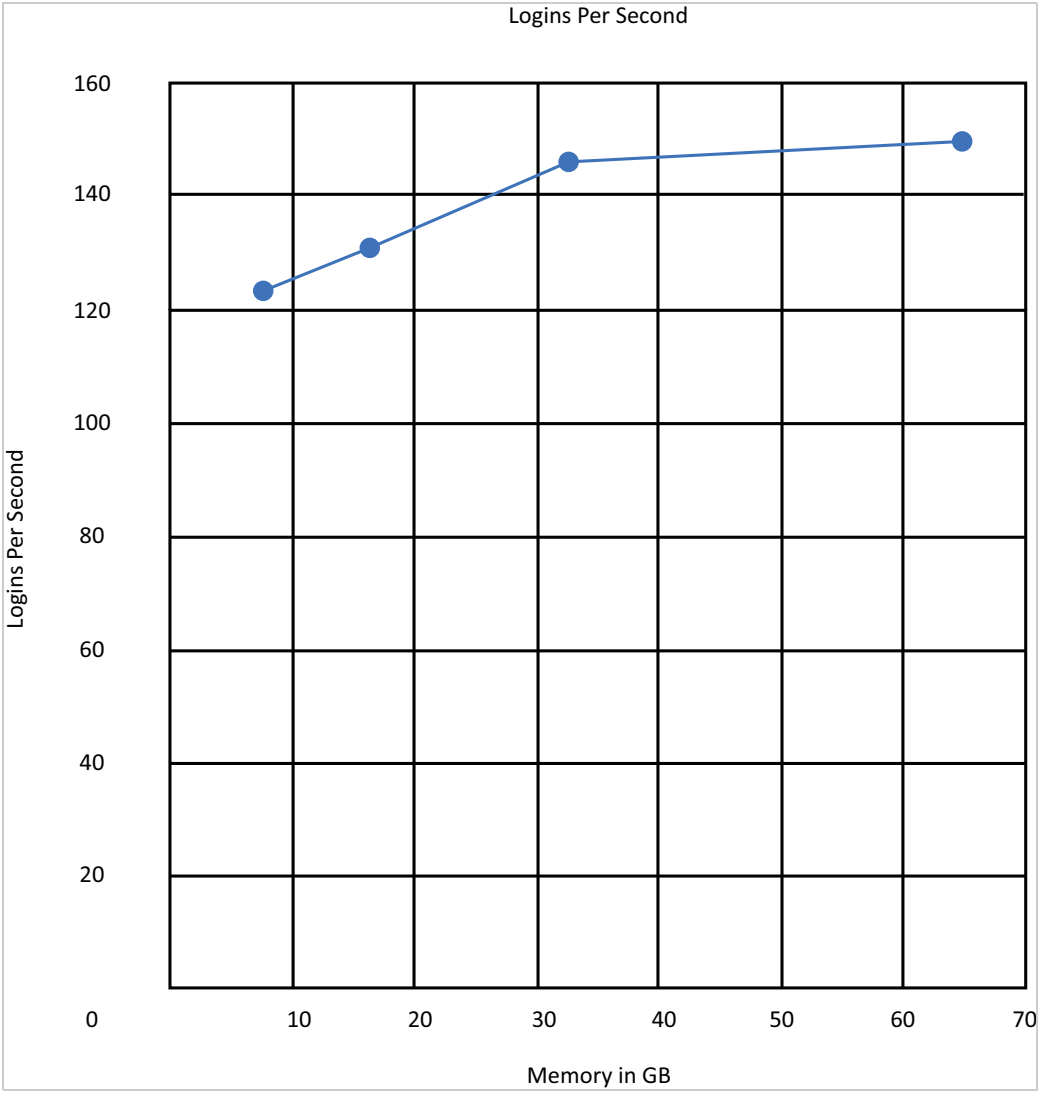

### Login Performance with Number of Nodes in the Cluster

In this test, each node is assigned 8 CPU and 16 GB memory. Performance is measured by increasing the number of nodes in the cluster.

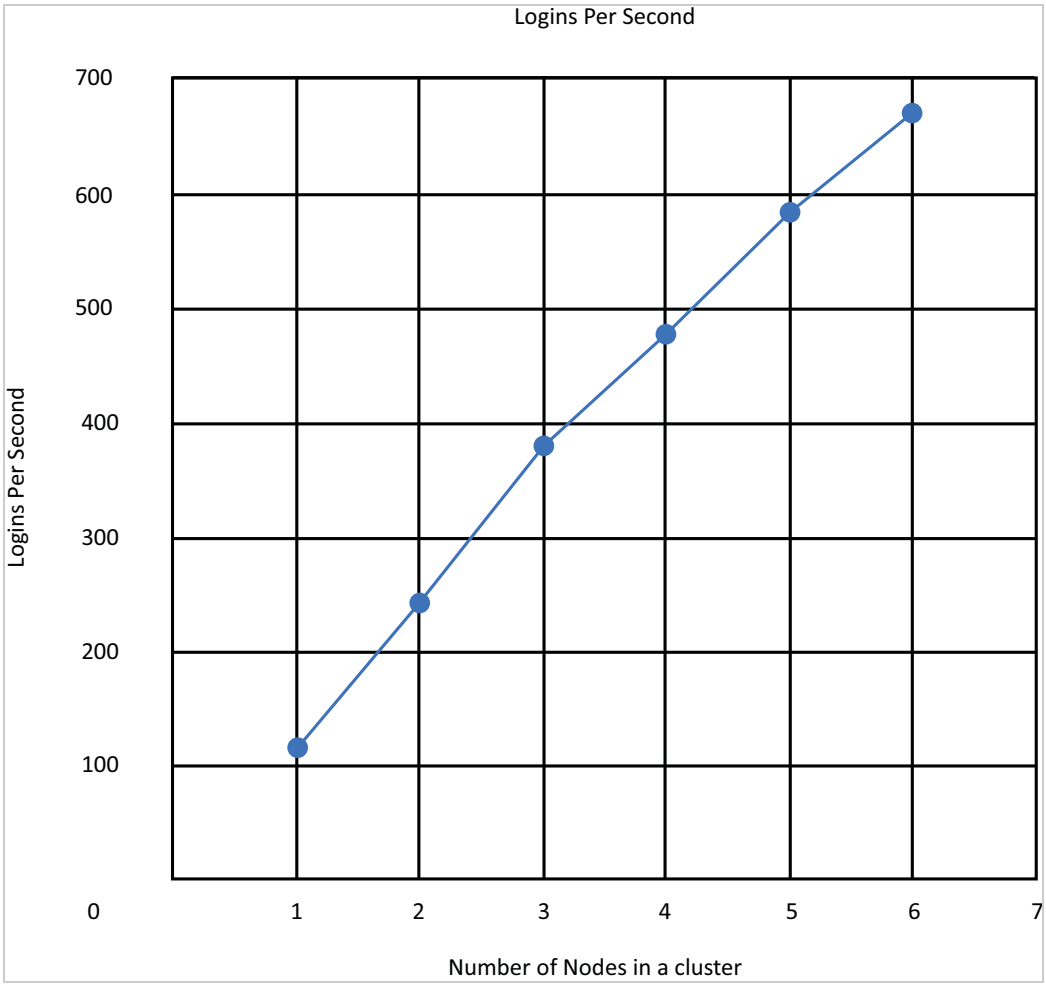


## 2.4.2 Scalability of Active Users Sessions

Test: Scaling and maintaining the active users sessions by periodically increasing users' logins and refreshing the active session within session timeout period.


- "Active Sessions Scalability with Scaling CPU" on page 18
- "Active Sessions Scalability with Scaling the Memory" on page 18

## Active Sessions Scalability with Scaling CPU

In this test, memory is kept constant at 32 GB and Tomcat is assigned 16 GB in Identity Server and Access Gateway. The number of CPU is increased in the order 1, 2, 4, 8, and 16, and performance is measured at each CPU level.

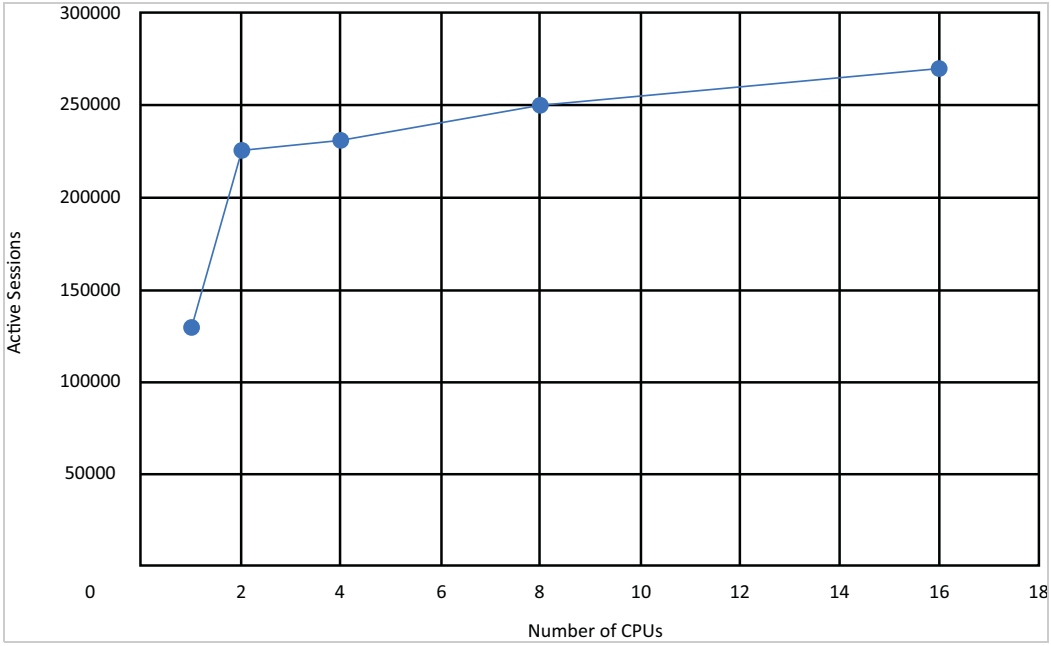


## Active Sessions Scalability with Scaling the Memory

In this test, the number of CPU is kept constant at 16 for Identity Server and Access Gateway. Memory is increased in the order 8 GB, 16 GB, 32 GB, and 64 GB. Tomcat is assigned with 70% of the available memory. Performance is measured at each memory level.

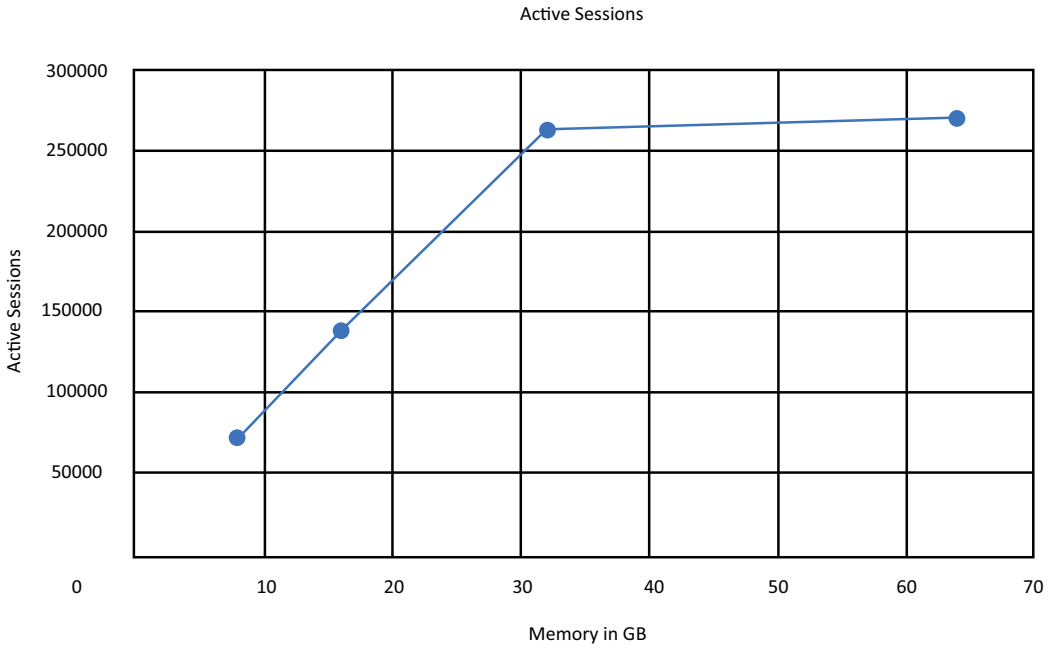

## 2.4.3 Access Gateway Hits Scalability

Test: Accessing a public resource through Access Gateway. Public resources are static pages of size 60 KB containing several hyperlinks to the same originating web server. In this test, the number of hits per second is measured.


- "Access Gateway Hits with Scaling CPU" on page 19
- "Access Gateway Hits with Scaling the Memory" on page 20


### Access Gateway Hits with Scaling CPU

The memory is kept constant to 32 GB and Tomcat is assigned 16 GB for Access Gateway. The number of CPU is increased in the order 1, 2, 4, 8, and 16. Performance is measured at each CPU.

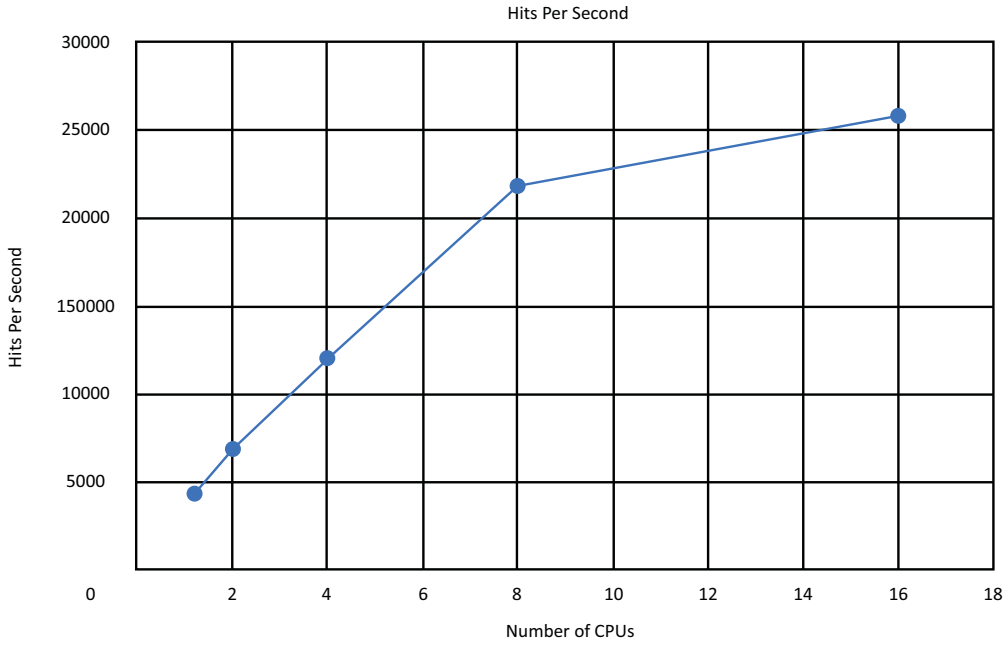

### Access Gateway Hits with Scaling the Memory

In this test, CPUs are kept constant at 16 for Access Gateway. Memory is increased in the order 8 GB, 16 GB, 32 GB, and 64 GB. Tomcat is assigned with 70% of the available memory. Performance is measured at each memory level.

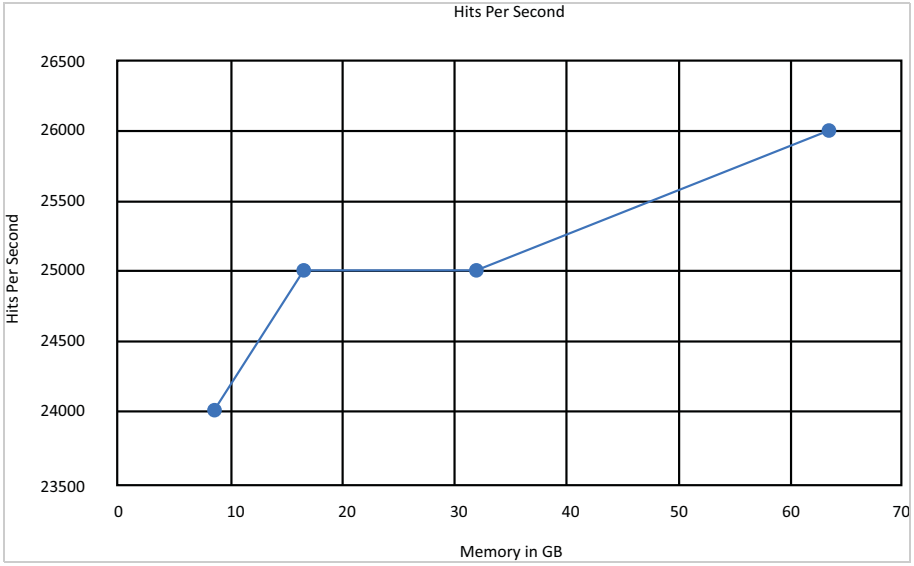


## 2.4.4 Access Gateway Throughput Scalability

Test: Accessing a public resource through Access Gateway. The public resources are static pages of size 8 MB that have several hyperlinks pointing to the same originating web server. Measure the throughput per second.

- “Access Gateway Throughput with Scaling CPU” on page 21
- “Access Gateway Throughput by scaling the Memory” on page 22

## Access Gateway Throughput with Scaling CPU

In this test, memory is kept constant at 32 GB and Tomcat is assigned 16 GB for Access Gateway. The number of CPU is increased in the order 1, 2, 4, 8, and 16. The performance is measured at each CPU level.

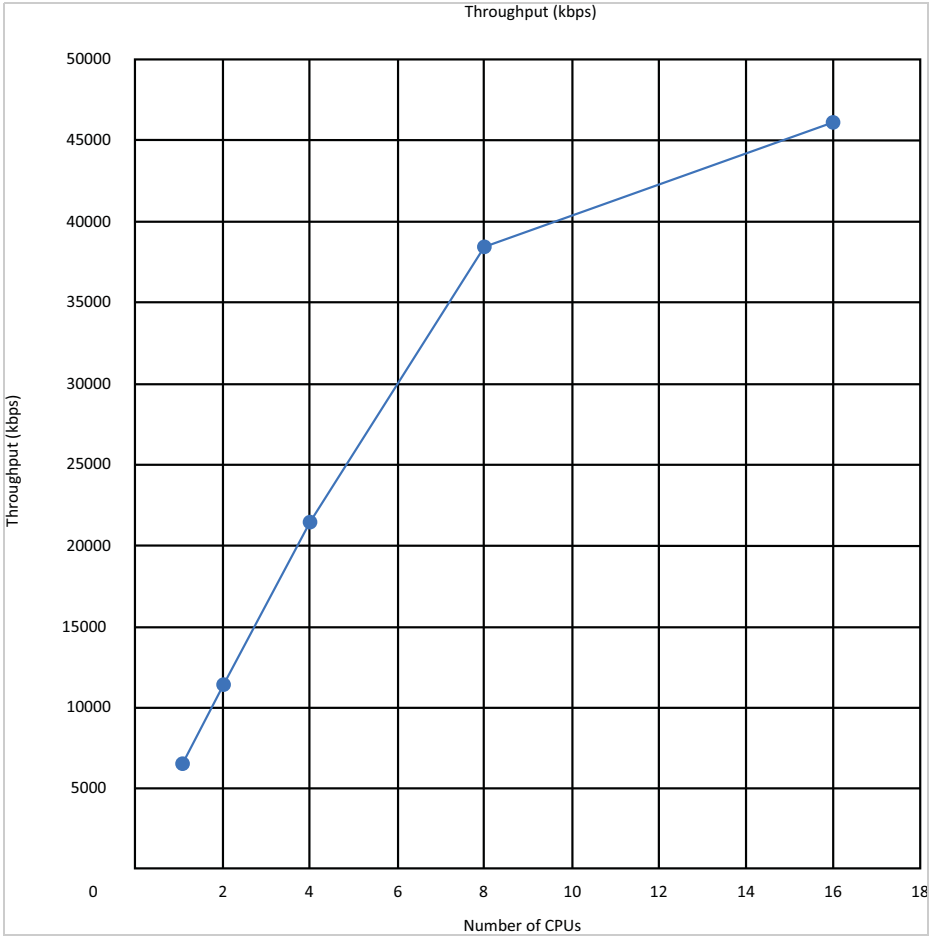

### Access Gateway Throughput by scaling the Memory

In this test, the number of CPU is kept constant at 16 for Access Gateway. Memory is increased in the order – 8 GB, 16 GB, 32 GB, and 64 GB. Tomcat is assigned with 70% of the available memory. Performance is measured at each memory level.

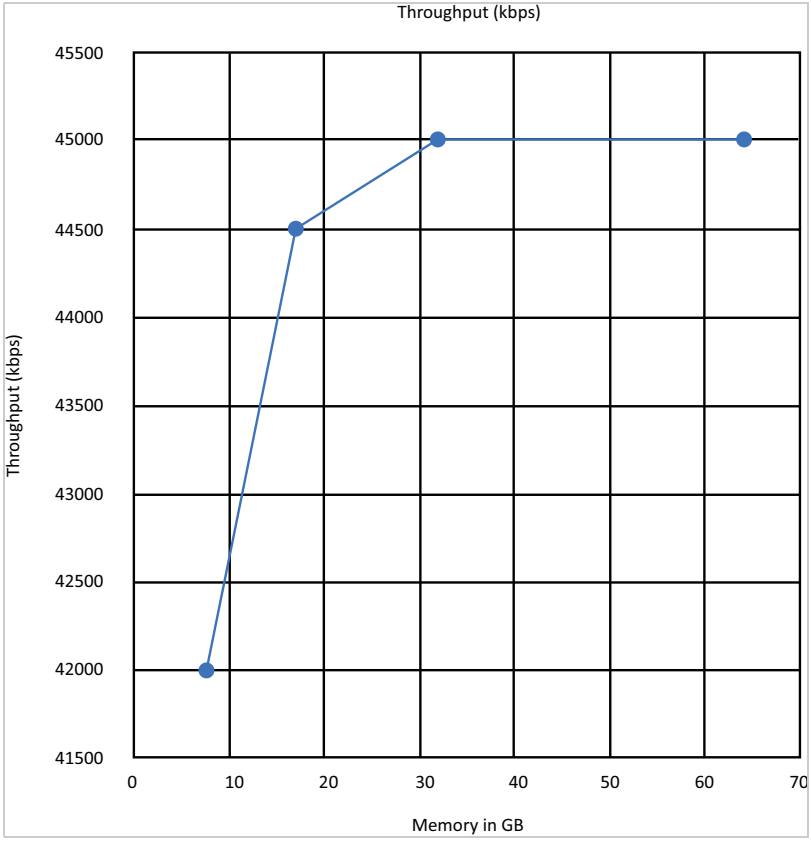


## 2.5 Sizing Recommendation for Analytics Server



### 2.5.1 Hardware Requirements for Analytics Server

For the demonstration purpose, the 50 GB hard disk is required. For a production environment, the hard disk requirement depends on the Access Manager login pattern for a day. For other system requirements for Analytics Server, see “System Requirements of Analytics Server”.

The following recommendations consider only Analytics Server-specific Access Manager Audit events. For information about Analytics Server events, see “Enabling Events for Each Graph” in the NetIQ Access Manager 5.0 Administration Guide.

Any change in Access Manager Audit events selection changes the disk requirement.

| | 25000 logins per day | 50000 logins per day | 100000 logins per day |
|---|---|---|---|
| **Number of days** | **Disk space in GB** | **Disk space in GB** | **Disk space in GB** |
| 1 | 0.058 | 0.115 | 0.23 |
| 10 | 0.575 | 1.15 | 2.3 |
| 30 | 1.725 | 3.45 | 6.9 |
| 60 | 3.45 | 6.9 | 13.8 |
| 90 | 5.175 | 10.35 | 20.7 |
| 120 | 6.9 | 13.8 | 27.6 |
| 180 | 20.99 | 41.975 | 83.95 |

### 2.5.2 Analytics Server Data Retention

The events stored in Analytics Server are retained in the local storage for 180 days. After 180 days, all events are purged from Analytics Server.

# 3 Access Gateway Performance in Access Manager 4.4

Access Manager 4.4 onward, Access Gateway is upgraded to Apache 2.4. Therefore, Access Gateway performance is significantly improved.

The following graphs show the overall public request performance improvement in Access Manager 4.4 over Access Manager 4.3:

**HTTPS transactions per second**

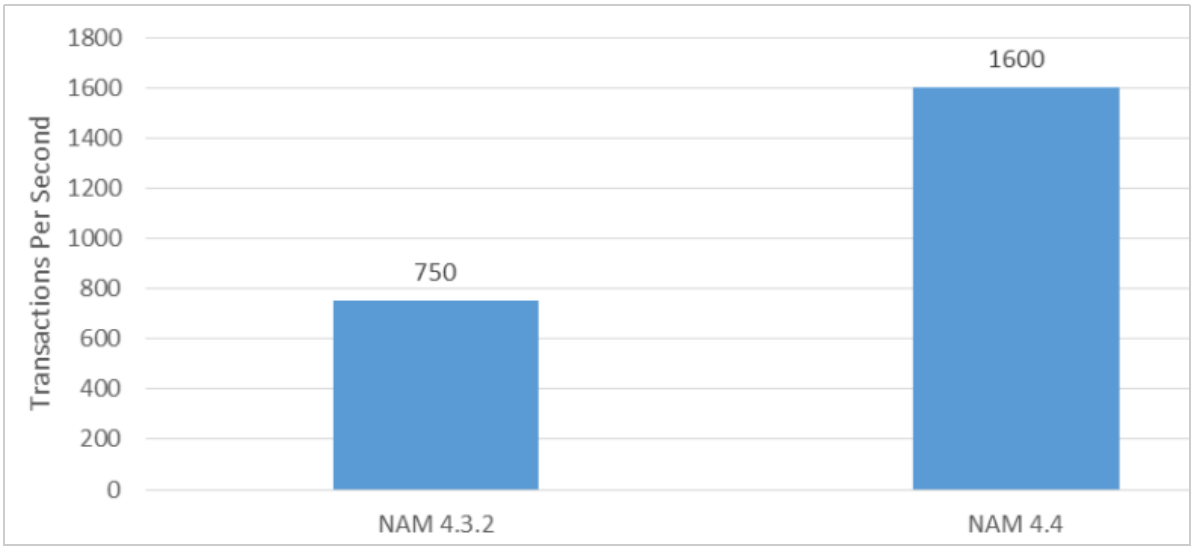


**HTTP transactions per second**

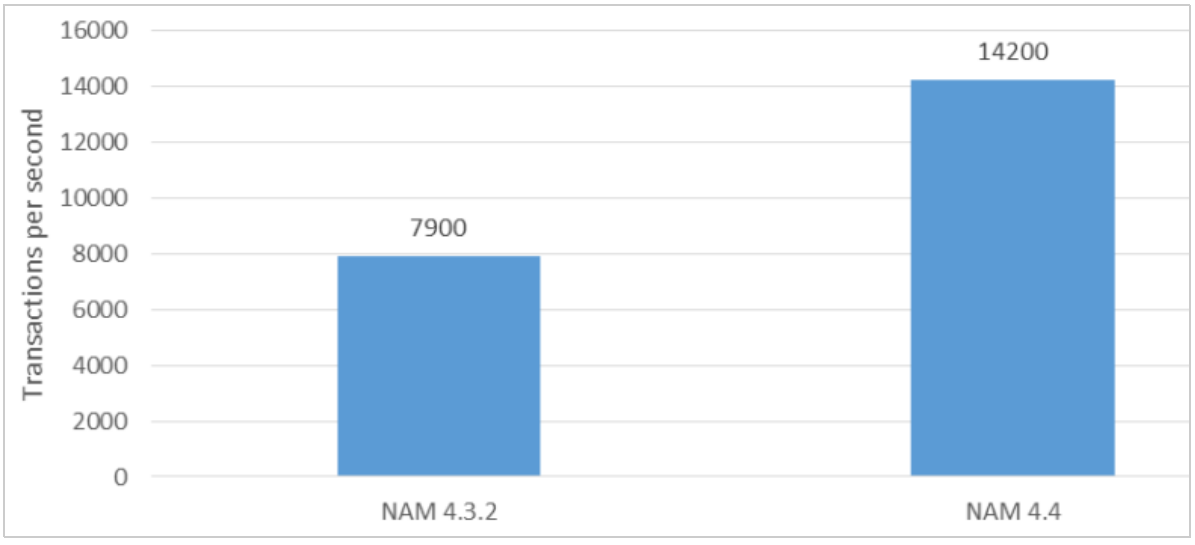

**Data transfer rate – HTTPS**

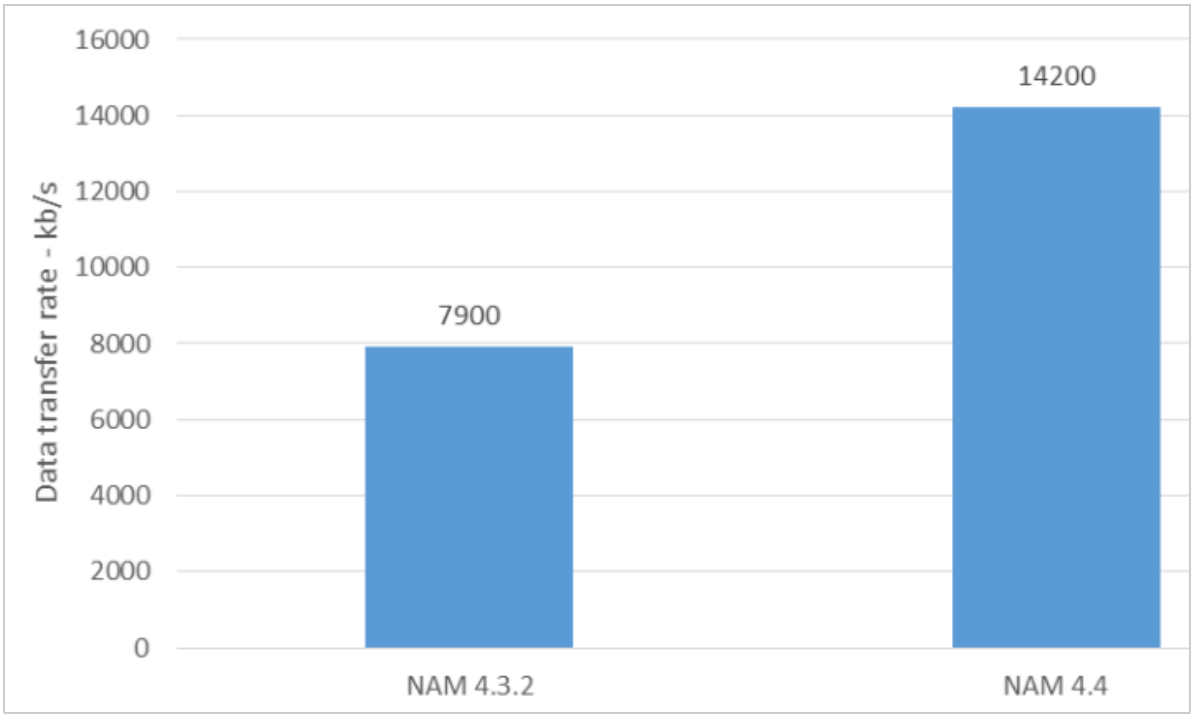


**Data transfer rate – HTTP**

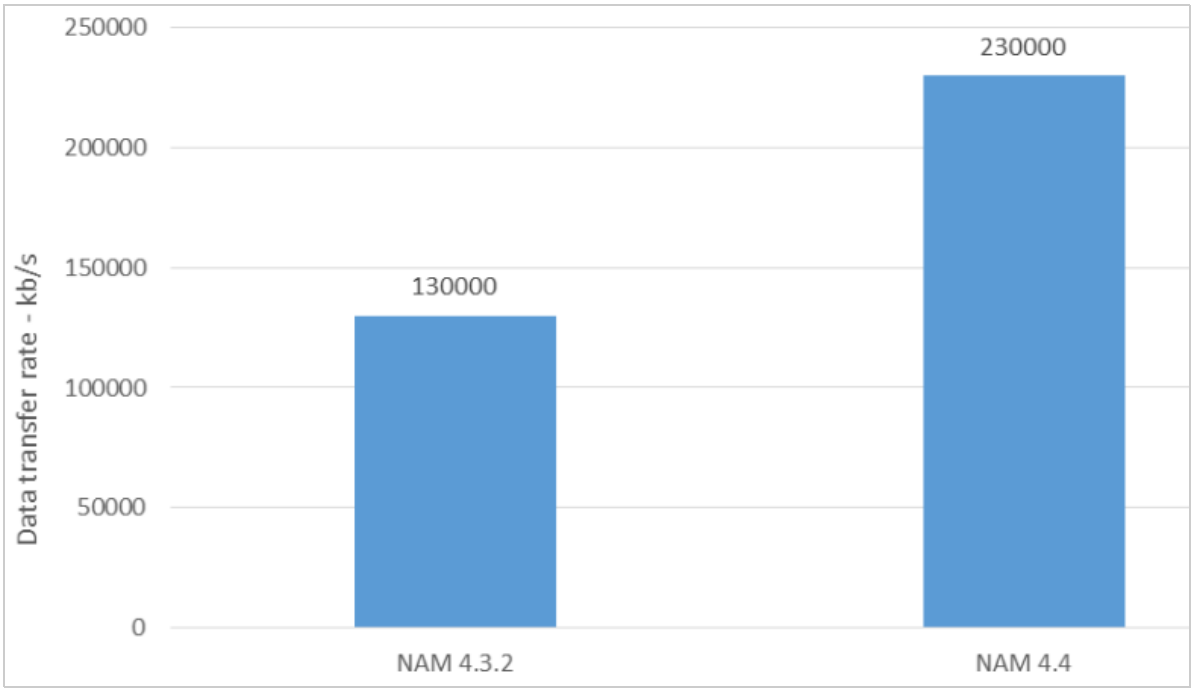

# A Additional Information



## A.1 Test Strategy

The test setup represents a medium-sized business with heavy traffic to help predict performance for both smaller and larger implementations. The performance, reliability, and scalability tests cover the critical areas that you need to know for designing your system.

A sizing guide is included to help determine the number of users that can be supported on a specific number of servers and configuration.

The tests cover the following major functional areas of public access, authentication, and authorization:

- The public requests test is focused on Access Gateway as a reverse proxy with caching to help increase the speed of your web servers by eliminating any authentication and authorization policy overhead.
- The authentication requests test is focused on the distributed architecture that provides a secure login to Access Manager.
- The authorization requests test is focused on the policy evaluation that occurs after the login has been completed and before the page is accessed.

The test environment includes a cluster of four Identity Servers and four Access Gateways. The number of users and the amount of traffic determine the size of the cluster.

**In this Section**



### A.1.1 Performance, Reliability, Scalability, and Failover Testing for Access Gateway

The performance testing includes the following scenarios:

- HTTP traffic through a public resource
- HTTPS traffic through a public resource
- HTTPS traffic through a protected resource

- HTTPS traffic through a protected resource with Form Fill
- HTTPS traffic through a protected resource with Identity Injection
- HTTPS traffic through a protected resource with policies that contain roles
- HTTPS traffic through a protected resource with 10 additional page requests

The reliability testing includes the *HTTPS traffic for 2 weeks through a stress test* scenario

The scalability (clustering) testing includes the following scenarios:

- 2 x 4 x 4 (2 Administration Console, 4 Identity Server, and 4 Linux Access Gateway)
- 2 x 4 x 4 (2 Administration Console, 4 Identity Server, and 4 Access Gateway Appliance)

The failover testing includes the *HTTP/HTTPS traffic continues after a component failover* scenario

## A.1.2 Test Setup



### Server Hardware for Access Gateway Appliance

Access Gateway clustered tests are run on an virtualized environment setup containing the following servers:

- Dell PowerEdge R730xd running ESXi 5.5
- Dell PowerEdge R720xd running ESXi 5.5
- Dell PowerEdge R710 running ESXi 5.5
- Dell PowerEdge R710 running ESXi 3.5

The design of the virtual machine is as follows:

| Server Components | Operating System | Hardware |
|---|---|---|
| Administration Console (2 nodes) | SLES11 SP3 | CPU: 2 x 3 GHz and Memory: 4 GB |
| Identity Servers (4 nodes) | SLES11 SP3 | CPU: 4 x 3 GHz and Memory: 16 GB |
| Access Gateway Appliance (4 nodes) | SLES11 SP3 | CPU: 4 x 2.6 GHz and Memory: 16 GB |
| External eDirectory user store (3 nodes) | SLES11 SP1 | CPU: 2 x 3 GHz and Memory: 4 GB |
| Apache2 Web Server (3 nodes) | SLES11 SP1 | CPU: 2 x 3 GHz and Memory: 4 GB |

## Server Hardware for Access Gateway Service on SLES 12

The Access Gateway clustered tests are run on an virtualized environment setup containing the following servers:

- Dell PowerEdge R730xd running ESXi 5.5
- Dell PowerEdge R720xd running ESXi 5.5
- Dell PowerEdge R710 running ESXi 5.5
- Dell PowerEdge R710 running ESXi 3.5

The design of the virtual machine is as follows:

| Server Components | Operating System | Hardware |
|---|---|---|
| Administration Console (2 nodes) | SLES 12 | CPU: 2 x 3 GHz and Memory: 4 GB |
| Identity Servers (4 nodes) | SLES12 | CPU: 4 x 3 GHz and Memory: 16 GB |
| Access Gateway Service (4 nodes) | SLES12 | CPU: 4 x 2.6 GHz and Memory: 16 GB |
| External eDirectory user store (3 nodes) | SLES11 SP1 | CPU: 2 x 3 GHz and Memory: 4 GB |
| Apache2 Web Server (3 nodes) | SLES11 SP1 | CPU: 2 x 3 GHz and Memory: 4 GB |

**NOTE:** In this performance testing, Access Gateway is installed on SLES 12 servers with BTRFS as a file system. Identity Server is installed on SLES 12 with EXT3 as a file system (upgraded from SLES 11 SP3 to SLES12)

## Server Hardware for Access Manager Appliance

Tests are run on a virtualized environment setup containing the following servers:

- Dell PowerEdge R730xd running ESXi 5.5
- Dell PowerEdge R720xd running ESXi 5.5
- Dell PowerEdge R710 running ESXi 5.5
- Dell PowerEdge R710 running ESXi 3.5

The design of the virtual machine is as follows:

| Server Components | Operating System | Hardware |
|---|---|---|
| Access Manager Appliance (4 nodes) | SLES11 SP3 | CPU: 8 x 3 GHz and Memory: 32 GB |
| External eDirectory user store (3 nodes) | SLES11 SP1 | CPU: 2 x 3 GHz and Memory: 4 GB |
| Apache2 Web Server (3 nodes) | SLES11 SP1 | CPU: 2 x 3 GHz and Memory: 4 GB |

## Load Balancers

The following L4 switches are used as load balancers for the testing:

- Zeus ZXTM LB (software L4 switch)

- Brocade ServerIron ADX 1000 (hardware L4 switch)
- Alteon 3408 (hardware L4 switch)

### Configuration Details

- HTML pages are of approximately 50 KB with 50 small images embedded for all public page tests.
- A small HTML page of 200B with one hyperlink is used for authentication, authorization, identity injection, and form fill performance tests. These tests do not cover the page rendering performance.
- Access Manager user stores configuration contains 20 threads with 100,000 users in a single container. Multiple containers received the same performance, however these tests have been conducted with optimization and fast hardware. If you do not optimize and increase the speed of your hardware, performance will decrease. The primary user store used in the tests is eDirectory 8.8.6.

### Performance, Reliability, and Stress Tools

The HP Mercury LoadRunner tool is used for Identity Server and Access Gateway testing. This tool correctly replicates large IP ranges between multiple clients in a clustered environment. This allowed the tests to simulate real-world environments more closely with a real browser interaction with Internet Explorer and Firefox.

The following are the specifications of the LoadRunner tests:

- The virtual user has 500 threads among 17 clients. This is the optimal amount of threads before the system started to receive excessive login times.
- HTML-based scripts describing user actions have been used. This is listed under the recording level and the HTML advanced option. This type of script helps to clear the cached data inside the script and downloads all data linked to the page.

If you do not have a sufficient IP address setup for LoadRunner, you must use solid load balancing on the Layer 4 switch. You must have parameters for the users so that you do not use the same user for every connection.

## A.1.3 Other Factors Influencing Performance Information

In addition to the hardware and the test configuration described in the previous sections, the following factors in a network also affect the overall performance:

- **Customized Login Pages:** Login and landing pages play an important role in the user experience while accessing the resources protected by Access Manager. Consider the performance aspect of the page loading and rendering while developing the custom login JSP pages.
- **L4 Switches:** A slow or incorrectly configured switch can severely affect performance. System test recommends to plug clustered Access Manager components directly into the switch or to segment accordingly. Enabling sticky bit/persistence on the L4 switch is also important. When sticky bit/persistence is not enabled, the product can handle the traffic correctly, however may become slower up to 50%.
- **Network Bandwidth:** Gigabit copper networking is used throughout the testing process. This is a requirement for the product to meet the testing results. If you are running at 100 MB or have a slow Internet connection, the product cannot solve this Performance bottleneck.

- **Web Servers:** The application servers are a major cause for slowness because they process most of the information. The tests used static and dynamic pages with more than 50 images. The tests are based on the real-world traffic to give a general idea of response times less than one second. The public requests can vary widely based upon the size of the page, caching settings, and content.
- **LDAP User Stores:** This component can cause slowness depending upon configuration, hardware, and the layout of the directory. The user store is the most common problem with performance. Therefore, testing must be done with the LDAP user stores that is used in the environment. Expect adjustments if you are attempting to get the maximum speed out of the cluster for the different LDAP user stores. eDirectory is primarily used throughout the testing to give a baseline for the product.
- **Timeout:** If you run a performance test, you must factor in sessions that are stored on the server. The tests have a 5 minute timeouts so that the tests do not overrun the total users on the system of 100,000 active sessions on the cluster. You must consider this while planning for capacity testing on a cluster. Configuring the session timeout for a resource is dependent on the security requirement. If security is not the concern, the following are some of the recommendations to fine-tune the session timeout configuration to reap the best performance:
  - If users access a protected resource for a short duration and leave the session idle after accessing few pages, configuring a short session timeout for such resources is recommended. This enables the system to remove idle sessions faster from the system.
  - If users access a protected resource for a long duration, configuring a long session timeout is recommended. It reduces the internal traffic to update the user access and improve the overall performance of the system.
- **Users:** Ensure that you have enough users on the system to run the performance test. If you run 50 threads of logins against Access Manager with each one using the same user to authenticate, Access Manager matches each user and handles all 50 sessions as the sessions of one user. This skews test goals and results, because it is not a valid user scenario and invalidate the test results.

## A.2 Tuning Parameters

This section provides details of the parameters tuned during the performance test to optimize the system performance. You must configure these parameters based on your environments.

It is recommended to test these parameters in the staging environment before running in the production environment.


- Section A.2.1, “Tuning Identity Server Parameters,” on page 31
- Section A.2.2, “Tuning Access Gateway Parameters,” on page 33
- Section A.2.3, “Web Socket Scalability,” on page 35


### A.2.1 Tuning Identity Server Parameters


- “Tomcat Connector Maximum Thread Setting” on page 32
- “JAVA Memory Allocations” on page 32
- “LDAP Load Threshold Configuration” on page 33

## Tomcat Connector Maximum Thread Setting

This parameter enables Identity Server to handle more threads simultaneously to improve the performance. The thread number must be fine-tuned for every customer environment based on the number of attributes attached to a user session. When each user session holds a large number of attributes, each user session requires more heap memory. The available stack memory reduces as a result. If the number of threads configured in this scenario is high, Tomcat tries to spawn more threads and fails due to non-availability of the stack memory. You must fine-tune the number of threads based on the attribute usage.

In Identity Server's server.xml file, set the value of maxThreads to 1000 for 8443 as follows:

```
<Connector NIDP_Name="connector" SSLEnabled="true" URIEncoding="utf-8"
acceptCount="100" address="x.x.x.x" ciphers="XX, XX ,XX, XX"
clientAuth="false" disableUploadTimeout="true" enableLookups="false"
keystoreFile="/opt/novell/devman/jcc/certs/idp/connector.keystore"
keystorePass="p2SnTyZPHn9qe66" maxThreads="1000" minSpareThreads="5"
port="8443" scheme="https" secure="true"
sslImplementationName="com.novell.nidp.common.util.net.server.NIDPSSLImple
mentation" sslProtocol="TLS"/>
```

For information about how to modify a file, see "Modifying Configurations" in the *NetIQ Access Manager 5.0 Administration Guide*.

---

**NOTE:** For Access Manager Appliance, the port number is 2443.

---

## JAVA Memory Allocations

The Tomcat configuration file controls the amount of memory that Tomcat can allocate for Java.

If you have installed Identity Server on a machine with a minimum 4 GB memory, you can modify the tomcat.conf file to improve performance under heavy load as follows:

In Identity Server's tomcat.conf, set the following parameters:

- Replace the Xms and Xmx values to 2048: `JAVA_OPTS="-server -Xms2048m -Xmx2048m –Xss256k "`

  This enables the Tomcat process to come up with 2 GB pre-allocated memory. If your Identity Server machine has more than 4 GB memory, the recommendation is to allocate 50% to 75% of the memory to Identity Server Tomcat. This needs to be fine-tuned based on each customer's environment.
- Set Identity Server Tomcat to 12288 for both Xms and Xmx.
- Change the -Dnids.freemem.threshold value from 0 to a value between 5 and 15. This parameter prevents user sessions from consuming all memory and ensures that free memory is available for other internal Java processes to run. When this threshold is reached, the user receives a 503 server busy message and a threshold error message is logged to the catalina.out file.

  ```
  JAVA_OPTS="${JAVA_OPTS} -Dnids.freemem.threshold=10"
  ```

For information about how to modify a file, see "Modifying Configurations" in the *NetIQ Access Manager 5.0 Administration Guide*.

### LDAP Load Threshold Configuration

In Identity Server's web.xml, set ldapLoadThreshold to 600.

```
<context-param>
     <param-name>ldapLoadThreshold</param-name>
     <param-value>600</param-value>
</context-param>
```

This enables Identity Server to make connections to the LDAP user store up to 600.

## A.2.2 Tuning Access Gateway Parameters

You can configure the following settings to optimize the performance:

- "AJP Connector Maximum Thread Setting" on page 33
- "JAVA Memory Allocations" on page 33
- "Access Gateway Appliance Advanced Options" on page 34
- "Apache MPM Settings" on page 34

### AJP Connector Maximum Thread Setting

In Access Gateway's server.xml, set maxThreads="1000" for the port 9009 connector.

For information about how to modify a file, see "Modifying Configurations" in the *NetIQ Access Manager 5.0 Administration Guide*.

This parameter enables Access Gateway Appliance ESP to handle more threads simultaneously to improve the performance. The thread number needs to be fine-tuned for every customer environment based on the number of attributes attached to a user session. When each user session holds a large number of attributes, each user session needs more heap memory. The available stack memory reduces as a result. If a number of threads configured in this scenario is high, Tomcat tries to spawn more threads and fails due to non-availability of the stack memory. You need to fine-tune the number of threads based on the attribute usage.

### JAVA Memory Allocations

The Tomcat configuration file controls the amount of memory that Tomcat can allocate for Java.

If you have installed Access Gateway on a machine with the minimum 4 GB of memory, you can modify the tomcat.conf file to improve performance under heavy load as follows:

- In Access Gateway's tomcat.conf, replace values of Xms and Xmx to 2048: JAVA_OPTS="-server -Xms2048m -Xmx2048m –Xss256k "

  This enables the Tomcat process to come up with 2 GB pre-allocated memory.

  For information about how to modify a file, see "Modifying Configurations" in the *NetIQ Access Manager 5.0 Administration Guide*.
- If the Access Gateway Appliance machine has more than 4 GB memory, the recommendation is to allocate 50% to 75% of the memory to ESP Tomcat. This needs to be fine-tuned based on each customer environment.

- Set Xms and Xmx to 12288 for ESP Tomcat.
- Change the -Dnids.freemem.threshold value from 0 to a value between 5 and 15.This parameter prevents user sessions from using up all memory and ensures that free memory is available for other internal Java processes to function. When this threshold is reached, the user receives a 503 server busy message and a threshold error message is logged to the catalina.out file. JAVA_OPTS="${JAVA_OPTS} -Dnids.freemem.threshold=10"

## Access Gateway Appliance Advanced Options

Add the following advanced option:

```
NAGGlobalOptions ESP_Busy_Threshold=5000
```

## Apache MPM Settings

In httpd-mpm.conf, `mpm_worker_module` is configured by default with the following settings:

```
<IfModule mpm_worker_module>
     ThreadLimit 300
     StartServers 3
     MaxClients 3000
     MinSpareThreads 3000
     MaxSpareThreads 3000
     ThreadsPerChild 300
     ServerLimit 10
     MaxRequestsPerChild 0
</IfModule>
```

This configuration is for the Appliance machine with the minimum 4 GB memory. If the Appliance machine has more than 6 GB memory, set mpm_worker_module to match the following configuration.

For information about how to modify a file, see “Modifying Configurations” in the *NetIQ Access Manager 5.0 Administration Guide*.

The performance tests are conducted with the following configuration when the Appliance machine has 16 GB memory available:

```
<IfModule mpm_worker_module>
     ThreadLimit 1000
     StartServers 9
     ServerLimit 10
     MaxClients 9000
     MinSpareThreads 9000
     MaxSpareThreads 9000
     ThreadsPerChild 1000
     MaxRequestsPerChild 0
</IfModule>
```

If the available memory is less or more, you must fine-tune each of these configurations based on your environment.

### A.2.3 Web Socket Scalability

Access Manager 4.4 onward, Access Gateway supports web socket applications. The scalability of Access Gateway for web socket connections depends on the Access Gateway hardware configuration along with a proper system and Access Gateway tuning.

Consider the following tuning for web socket scalability:

- ***Maximum number of open files for Access Gateway:***

  In large scale Web-Socket deployments, Access Gateway may run out of the available maximum number of open file descriptor after reaching the default maximum open file descriptor. It is recommended to configure more number of open file descriptor in such cases. To find the maximum number of open files for a process, run the following command on the Linux server to know the maximum number of open files for the process:

  ```
  #ulimit -n
  ```

- ***Apache Multi Processing Modules (MPM) Tuning:***

  Access Gateway requires independent threads to handle individual Web-Socket requests. apache httpd-mpm.conf must be tuned properly based on the web socket traffic that is expected to load to the Access Gateway server. For example, using the following configuration, you can scale 30K Web-Socket connections on an Access Gateway node:

  - Hardware: 4 CPU, 16 GB Memory
  - Ulimit setting: ulimit -n 8192

- ***httpd-mpm.conf:***

  Edit the following setting in httpd-mpm.conf:

  ```
  <IfModule mpm_worker_module>
      ThreadLimit 3000
      StartServers 9
      ServerLimit 10
      MaxClients 30000
      MinSpareThreads 9000
      MaxSpareThreads 9000
      ThreadsPerChild 3000
      MaxRequestsPerChild 0
  </IfModule>
  ```

  For information about how to modify a file, see "Modifying Configurations" in the *NetIQ Access Manager 5.0 Administration Guide*.

## A.3 Test Environment: Identity Server as an OAuth 2.0 Identity Provider

Tests are run on Access Manager 4.3. Later versions of Access Manager are expected to behave in the similar manner.

- Server Hardware

- Access Manager Tuning
- Test Tools

### Server Hardware

The tests are run on a virtualized lab with the following configuration:

| Hardware | Virtual Machines |
|---|---|
| Dell PowerEdge R7720xd<br>◆ CPU: 12 Cores @ 2.9GHz –<br>◆ RAM: 96 GB<br>◆ VMWare ESXi 5.5 | Administration Console<br>◆ CPU: 4 Cores @ 2.9GHz<br>◆ RAM: 8 GB<br>LDAP User Store 1<br>◆ CPU: 4 Cores @ 2.9GHz<br>◆ RAM: 8 GB<br>LDAP User Store 2<br>◆ CPU: 4 Cores @ 2.9GHz<br>◆ RAM: 8 GB |
| Dell PowerEdge R710xd<br>◆ CPU: 12 Cores @ 2.9GHz<br>◆ RAM: 86 GB<br>◆ VMWare ESXi 5.5 | Identity Server<br>◆ CPU: 8 Cores @ 2.9GHz<br>◆ RAM: 16 GB |

### Access Manager Tuning

Customized login pages are used in the testing. The following settings are done in Identity Server:

- Tomcat is set with 8 GB memory in tomcat.conf

  ```
  JAVA_OPTS="-server -Xmx8192m -Xms8192m -Xss256k "
  ```
- Tomcat connector maximum threads is set to 1000 (maxThreads="1000") in `server.xml` for port 8443.
- LDAP load threshold (ldapLoadThreshold) is set to 600 in web.xml

### Test Tools

Silk Performer 17.0

## A.4 Test Environment: Advanced Session Assurance

These tests are run on Access Manager 4.3. Later versions of Access Manager are expected to behave in the similar manner.

- Hardware Configuration
- Access Manager Tuning

- Test Tool
- Session Assurance Parameters

## A.4.1 Hardware Configuration

The hardware layout and distribution of virtual machines are only for the test purpose. It is recommended to not combine the critical resources on the same ESX server.

The tests are run on a virtualized lab with the following configuration:

| Hardware | Virtual Machines |
|---|---|
| Dell PowerEdge R7720xd<br>• CPU: 12 Cores @ 2.9GHz –<br>• RAM: 96 GB<br>• VMWare ESXi 5.5 | **Administration Console**<br>• CPU: 4 Cores @ 2.9GHz<br>• RAM: 8 GB<br>**LDAP User Store-1**<br>• CPU: 4 Cores @ 2.9GHz<br>• RAM: 8 GB<br>**LDAP User Store-2**<br>• CPU: 4 Cores @ 2.9GHz<br>• RAM: 8 GB |
| Dell PowerEdge R710xd<br>• CPU: 12 Cores @ 2.9GHz<br>• RAM: 86 GB<br>• VMWare ESXi 5.5 | **Identity Server**<br>• CPU: 8 Cores @ 2.9GHz<br>• RAM: 16 GB |
| Dell PowerEdge R730xd<br>• CPU: 16 Cores @ 2.9GHz –<br>• RAM: 132 GB<br>• VMWare ESXi 5.5 | **Access Gateway**<br>• CPU: 8 Cores @ 2.9GHz<br>• RAM: 16 GB<br>**Web Server-1**<br>• CPU: 4 Cores @ 2.9GHz<br>• RAM: 8 GB<br>**Web Server-2**<br>• CPU: 4 Cores @ 2.9GHz<br>• RAM: 8 GB |

## A.4.2 Access Manager Tuning

Login pages used in testing are customized. The following are component-specific settings:

### Identity Server:

- Tomcat is set with 8 GB memory in tomcat.conf.

```
JAVA_OPTS="-server -Xmx8192m -Xms8192m -Xss256k "
```

- Tomcat connector maximum threads is set to 1000 (maxThreads="1000") in server.xml for port 8443.
- LDAP load threshold (ldapLoadThreshold) is set to 600 in web.xml.

**Access Gateway:**

- Tomcat is set with 8 GB memory in tomcat.conf.

```
JAVA_OPTS="-server -Xmx8192m -Xms8192m -Xss256k "
```

- Tomcat connector maximum threads is set to 1000 (maxThreads="1000") in server.xml for port 8443.
- Apache mpm_worker_module in httpd-mpm.conf is configured with the following settings:

```
<IfModule mpm_worker_module>
    ThreadLimit 300
    StartServers 3
    MaxClients 3000
    MinSpareThreads 3000
    MaxSpareThreads 3000
    ThreadsPerChild 300
    ServerLimit 10
    MaxRequestsPerChild 0
</IfModule>
```

### A.4.3 Test Tool

Silk Performer 17.0 is used in testing.

### A.4.4 Session Assurance Parameters

The performance tests are run by enabling the following default session assurance parameters:

- Request Header Set
- Hardware Parameters
- Language Set
- Operating System Parameters
- TimeZone Offset
- User Agent

## A.5 Test Environment: Vertical and Horizontal Scalability

In vertical scaling, the capacity (CPU and memory) of a single instance of an Access Manager component is increased and tested for performance and scalability.

In horizontal scaling, additional instances of Access Manager components are added to the component cluster and tested for performance and scalability.

- Test Infrastructures
- Test Configuration and Test Data
- Access Manager Tuning

## A.5.1 Test Infrastructures

The test lab consists of virtualized isolated environment where in test servers are running as virtual machines on top of the VMWare ESXi Server.

The following diagram illustrates the virtual machine layout of a **Vertical Scaling** setup:

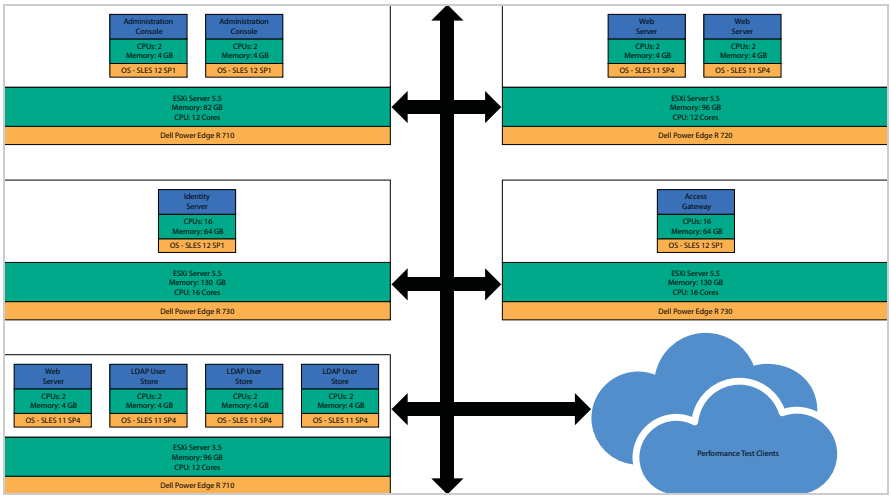


The following diagram illustrates the virtual machine layout of a **Horizontal Scaling** setup:

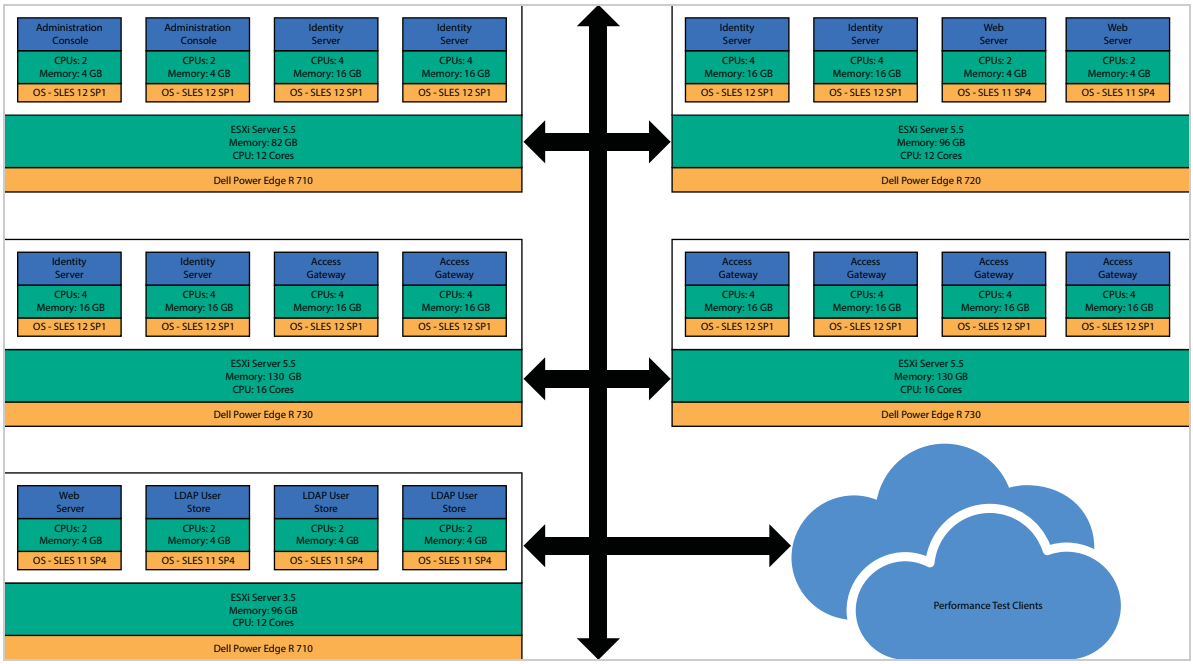


## A.5.2 Test Configuration and Test Data

- For login performance tests, the eDirectory user store with 3 replicas is used. These replicas have 100,000 users, which are synced across all replicas. A Secure Name Password Form authentication contract is used with a session time out of 10 minutes.

- For active sessions scaling tests, the eDirectory user store with a single replica having 1,000,000 users is used. A Secure Name Password Form authentication contract is used with a session time out of 30 minutes.
- For Access Gateway throughput and hit tests, 3 web servers with the static web pages of size 60 KB are used.
- The performance tests are run with Borland Silk Performer version 16.5.
- During the tests, the load test clients and the servers were using the TSL_DHE_RSA_WITH_AES_128_CBC_SHA ciphers for SSL negotiation. Any change in the cipher may impact the performance behavior of Access Manager components.

## A.5.3 Access Manager Tuning

For detailed information about Access Manager tuning, see Tuning Parameters. The following details are specific to the vertical and horizontal scalability tests:


- "Identity Server Tuning" on page 40
- "Access Gateway Tuning" on page 40


### Identity Server Tuning

#### JAVA Memory Allocation

- Vertical memory scaling tests: 70% of the total memory is allocated to Tomcat.
- Horizontals scaling tests: 8 GB memory is allocated to Tomcat in each Identity Server configuration.

You can allocate the Java memory by modifying –Xms and Xmx values in tomcat.conf.

For information about how to modify a file, see "Modifying Configurations" in the *NetIQ Access Manager 5.0 Administration Guide*.

#### Tomcat Max open files

The value of Tomcat max open files is set to 16384. This can be set by adding the following entry in the tomcat.conf file:

```
ulimit -Hn 16384
ulimit -Sn 16384
```

For information about how to modify a file, see "Modifying Configurations" in the *NetIQ Access Manager 5.0 Administration Guide*.

### Access Gateway Tuning

You can allocate the Java memory by modifying –Xms and Xmx values in tomcat.conf.

- In vertical memory scaling tests, 70% of the total memory is given to Tomcat.
- In horizontals scaling tests, 8 GB memory is allocated to Tomcat in each Access Gateway configuration.

For information about how to modify a file, see "Modifying Configurations" in the *NetIQ Access Manager 5.0 Administration Guide*.