

ALM Octane 12.60.21 Benchmark Test Results (MS SQL Server)

Contents

Contents.....	2
Introduction	3
Test result recommendations	3
Test details.....	4
Test description.....	5
Data load details.....	5
Environment details	6
Test results.....	7
Two-node ALM Octane.....	7
Four-node ALM Octane	9
Appendix A – CentOS system changes before starting	11
Learn more.....	12
Micro Focus Trademark Information.....	12

Introduction

This document describes how resources are consumed in systems running ALM Octane version 12.60.21. The data provided in this document was obtained by running identical scale tests on two different-sized environments. This should help you plan the ALM Octane deployment within your organization. Note that results may vary depending on your network configuration, company policy, anti-virus software, and so on.

Test result recommendations

The following table presents a summary of the scale test results, as well as the recommended system configuration for each system size. (This is based on the number of **concurrent** users, which as a rule of thumb is about 10% of listed users).

No. of Concurrent users	0-200	200-599	600-1000
Application server configuration	2 nodes, each 4 CPU, 8 GB RAM, Max 4 GB heap	2 nodes, each 8 CPU, 12 GB RAM, Max 8 GB heap	4 nodes, each 8 CPU, 12 GB RAM, Max 8 GB heap
DB server configuration	8 CPU, 16 GB RAM	12 CPU, 32 GB RAM	24 CPU, 64 GB RAM
Elasticsearch node (3 nodes min)	2 CPU (and up), 6 GB RAM	4 CPU (and up), 8 GB RAM	8 CPU (and up), 16 GB RAM

When estimating sizes for your system, experiment with different CPU sizes depending on the number of concurrent users you require.

Test details

The scale test we performed is an extreme case based on synthetic data, which is created based on real customers' data volume and flows.

This section includes the following:

- [Virtual user types](#)
- [Test description](#)

Virtual user types

The types of virtual users used in the scale test include the following:

- [A – Viewing virtual user \(78% of total actions\)](#)
- [B – Add/Update items virtual user \(17% of total actions\)](#)
- [C – Login virtual user \(5% of total actions\)](#)

Percentage of transactions

Percentage of time	Action
32%	Filter tables
17%	View entity tables
16%	View Dashboards
10%	Create new entity (15% of the 68% perform it as bulk operation)
10%	Search entities
5%	Log into ALM Octane, wait 10 seconds, and then log out. Upon login, the last view opens.
5%	View entity details
4%	Update existing entity
1%	Delete existing entity

Test description

On average, every 75 seconds a new user is added to the system load. The test ends when the system exhausts its ability to support users. We then run for 24 hours with the maximum load.

Background scenario for receiving CI server data (ALM Octane plugin for Jenkins)

The following background processes run on 2 workspaces:

- 8 pipelines (total of 16 in the shared space, 32 for entire site).
- Each pipeline receives results of the pipeline run from Jenkins every 20 to 120 minutes.
- The pipeline run includes commits, related work items, related stack traces and 1000 runs.
- 10 pipelines are Dashboard viewers.

Data load details

The data load used in the test is determined based on real customer data on volume and flows. This includes the following:

10000 listed users

10 shared spaces: 2 large-scale shared spaces with the following:

50 workspaces with the following distribution:

- 5 workspaces (90% of users), where each workspace contains:
 - 5 releases
 - 400 products
 - 450 epics. Each epic contains an average of 5-6 features, totaling up to 2500 features for all releases.
 - 9500 user stories for all releases
 - 16000 defects per workspace
 - Tests – 4000 manual tests, 20000 manual runs, 5000 automatic tests, 40000 automated runs, and 3000 test suites.
 - 2 CI servers
 - 4 pipelines per CI server (8 in total)
 - 40000 CI builds
 - 219 business rules
- 20 workspaces, each containing 10% of the content of the above 5 workspaces. 10% of users will log in to these.
- The remaining 25 workspaces are almost empty. They function simply as background load, and no virtual users log in to them.

10 additional shared spaces with about 20% of large-scale shared space data load as described above.

Environment details

All environments are hosted on a dedicated virtual machine (other than the Database server).

No environments are shared, and all resources are reserved and isolated. Each virtual machine is a physical grade server, and includes the following types of hardware:

Hardware configuration

Component	Configuration
ALM Octane node	<ul style="list-style-type: none">• 8 cores CPU• 16 GB RAM• CENTOS 7.4 64 bit• ALM Octane Enterprise license
Database server	Physical HPE DL 380 G9 <ul style="list-style-type: none">• Dual Intel(R) Xeon(R) CPU E5-2640 v3 @ 2.60GHz (total 16 cores + HT)• 128 GB RAM• Raid 0 for OS, Raid 5 for DATA• SQL Server 12.0.4100.1
Elasticsearch servers	<ul style="list-style-type: none">• 2 nodes, each has 8 cores CPU, 16 GB RAM, max heap size set to 8 GB, Windows 2012 R2 64 bit
NFS file server	4 cores CPU, 4GB RAM, 350 GB NFS share, CENTOS 7.4 64 bit
Simple software load balancer	4 cores CPU, 4GB RAM, CENTOS 7.4 64 bit, Nginx 1.12.2

Test results

This section presents the test results for each of the configurations described above. This section includes:

[Two-node ALM Octane](#)

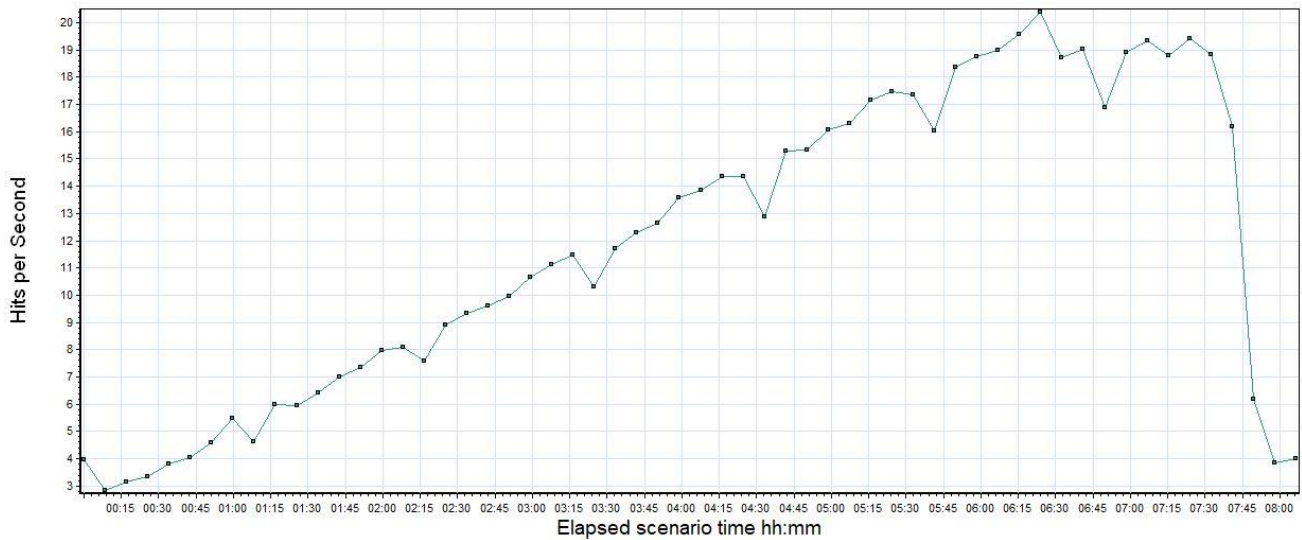
[Four-node ALM Octane](#)

Two-node ALM Octane


This configuration includes two clustered ALM Octane nodes with the same configuration described in the hardware section.

No. of Concurrent users	100	200	300	400	500
Single App server CPU utilization (max stable 700%)	150%	260%	335%	420%	480%
Used Java heap size (in GB, max 8 GB)	2.3 GB	2.5 GB	2.8 GB	3.1 GB	3.1 GB
SQL Server CPU utilization (max stable 3200%)	176%	290%	375%	557%	633%
Elasticsearch server CPU	<200%	<200%	<200%	<200%	<200%
Elasticsearch server Java heap size	<3	<3	<5	<5	<5
Hits per second (server load)	5	8	11	13.5	16

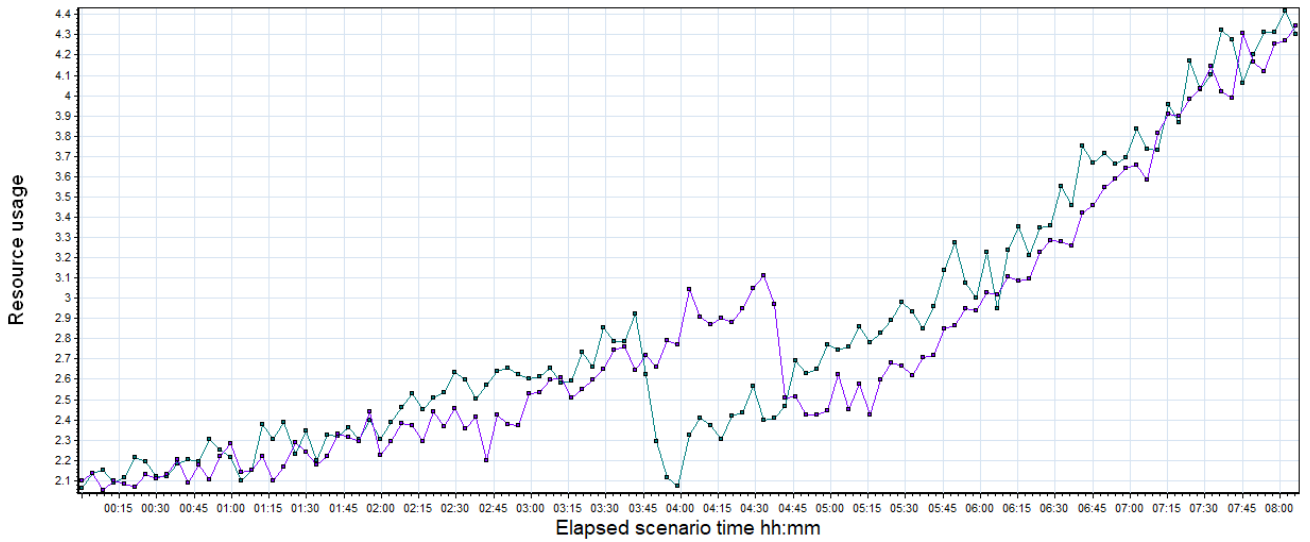
Hits per second



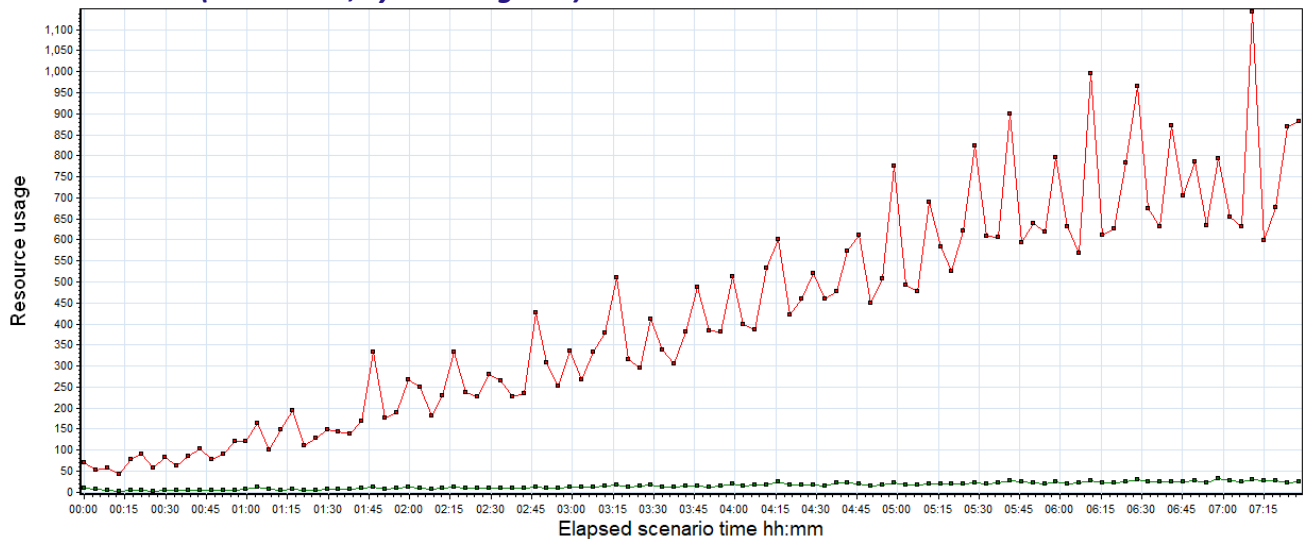
Application server CPU (user in red, system in green)



Used Java heap size (in GB, node 1 in green, node 2 in purple)



Database CPU (user in red, system in green)




Four-node ALM Octane

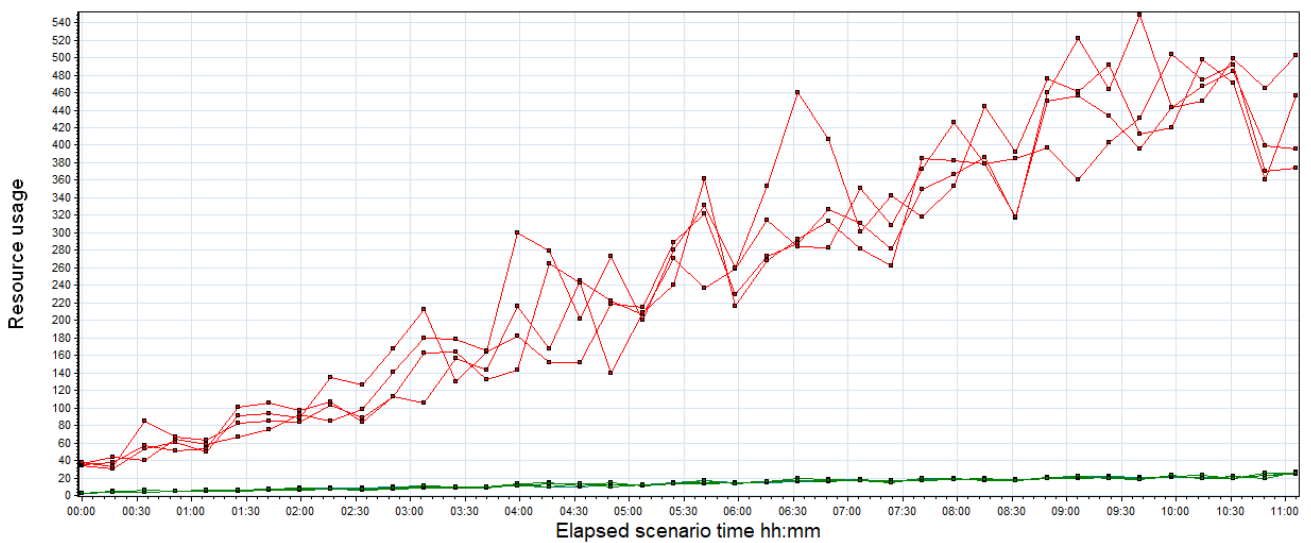
This configuration includes 4 clustered ALM Octane nodes with the same configuration described in the hardware section.

No. of Concurrent users	600	800	1000
Single App server CPU utilization (max stable 700%, min-max)	350%	440%	550%
Used Java heap size (in GB, max 12 GB)	3.2 GB	3.8 GB	5.5 GB
DB server CPU utilization (max stable 3200%)	832%	1170%	1408%
Elasticsearch server CPU	<400%	<400%	<400%
Elasticsearch server Java heap size (in GB)	<7	<7	<8
Hits per second (server load)	20	25	29

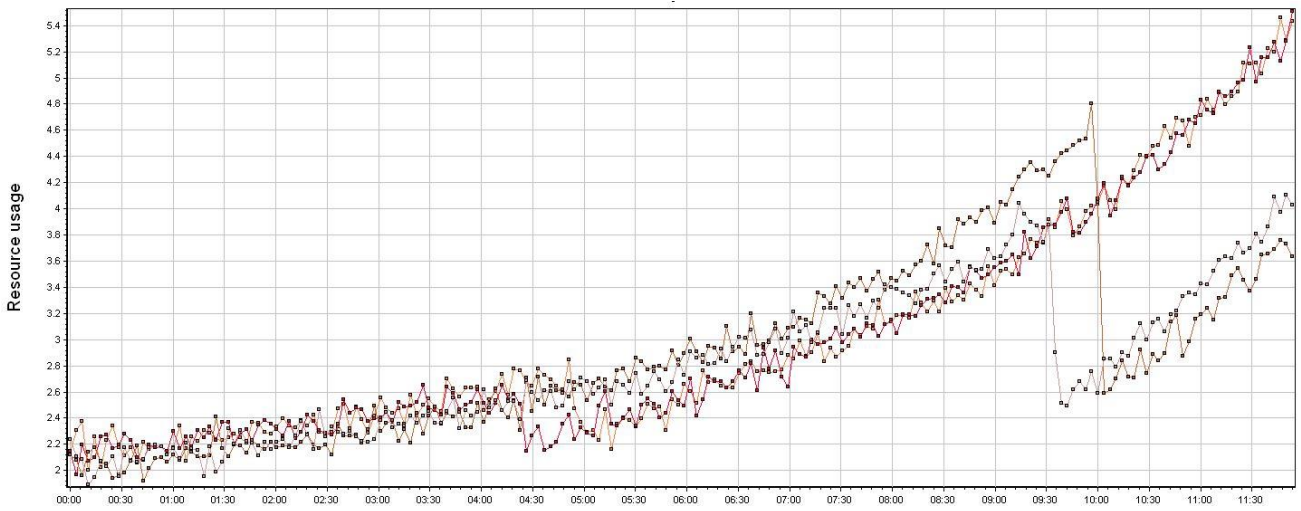
Hits per second





Application server CPU (4 nodes, user in red, system in green)



Used Java heap size (4 nodes)



Database CPU



Appendix A – CentOS system changes before starting

Before starting the test, we modified the maximum number of open files on the CentOS 7 system. The default maximum number is 1024. We modified this to a minimum of 8000 files for 100 users.

Increase this number by 1000 for every additional 100 concurrent users.

Consult your Linux distribution documentation on how to make this permanent. For CentOS & RHEL, change **/etc/security/limits.conf** as described in <https://access.redhat.com/solutions/61334>.

You can modify this setting at every restart, using the following command:

ulimit <limit>

For example, to set the limit to 8000, enter:

ulimit 8000

Learn more

www.microfocus.com/alm-octane

Micro Focus Trademark Information

MICRO FOCUS and the Micro Focus logo, among others, are trademarks or registered trademarks of Micro Focus (IP) Limited or its subsidiaries in the United Kingdom, United States and other countries. All other marks are the property of their respective owners.