

A decorative pattern of overlapping, multi-lined hexagons in a light blue color, set against a dark blue background, located at the top of the page.

FortiSOAR - Performance Benchmarking

Version 6.4.4



FORTINET DOCUMENT LIBRARY

<https://docs.fortinet.com>

FORTINET VIDEO GUIDE

<https://video.fortinet.com>

FORTINET BLOG

<https://blog.fortinet.com>

CUSTOMER SERVICE & SUPPORT

<https://support.fortinet.com>

FORTINET TRAINING & CERTIFICATION PROGRAM

<https://www.fortinet.com/support-and-training/training.html>

NSE INSTITUTE

<https://training.fortinet.com>

FORTIGUARD CENTER

<https://fortiguard.com/>

END USER LICENSE AGREEMENT

<https://www.fortinet.com/doc/legal/EULA.pdf>

FEEDBACK

Email: techdoc@fortinet.com

January, 2021

FortiSOAR 6.4.4 Performance Benchmarking

00-400-000000-20210112

TABLE OF CONTENTS

Change Log	4
FortiSOAR Performance Benchmarking for v6.4.4	5
Single Invocation Test for the single-node FortiSOAR appliance	5
Environment	5
Pre-test conditions on both the standalone FortiSOAR machine and the FortiSOAR High Availability (HA) cluster	6
Notes on the tests performed	6
Test setup for the single-node FortiSOAR appliance	6
Tests performed	6
Sustained Invocation Test for the single-node FortiSOAR appliance	9
Results	10
Graphs	10
Single Invocation Test for the High Availability (HA) active-active cluster of two FortiSOAR nodes	12
Test setup for the HA active-active cluster of two FortiSOAR nodes	12
Tests performed	12
Sustained Invocation Test for the HA active-active cluster of two FortiSOAR appliance	15
Results	16
Graphs	16

Change Log

Date	Change Description
2021-01-12	Initial release of 6.4.4

FortiSOAR Performance Benchmarking for v6.4.4

This document details the performance benchmark tests conducted in Fortinet labs. The performance benchmarking tests were performed on FortiSOAR version 6.4.4 Build 3164.

The objective of this performance test is to measure the time taken to create alerts in FortiSOAR, and complete the execution of corresponding playbooks on the created alerts in the following cases:

- Single-node FortiSOAR appliance
- Cluster setup of FortiSOAR

The data from this benchmark test can help you in determining your scaling requirements for a FortiSOAR instance to handle the expected workload in your environment.

Single Invocation Test for the single-node FortiSOAR appliance

Environment

FortiSOAR Virtual Appliance Specifications

Component	Specifications
CPU	8 CPUs
Memory	32 GB
Storage	250 GB virtual disk, with IOPS 2400, attached to an AWS Instance.

Operating System Specifications

Operating System	Kernel Version
CentOS 7	3.10.0-1160.6.1.el7.x86_64

External Tools Used

Tool Name	Version
Zabbix	4.2.1
Internal Script to gather data	

Pre-test conditions on both the standalone FortiSOAR machine and the FortiSOAR High Availability (HA) cluster

At the start of each test run -

- The test environment contained zero alerts.
- The test environment contained only the FortiSOAR built-in connectors such as IMAP, Utilities, etc.
- The system playbooks were deactivated such as, alert assignment notification, sla calculation, etc, and there were no running playbooks.
- The playbook execution logs were purged.
- Configured tunables as follows:
 - Changed celery workers to 16
 - Elastic heaps size to 8GB
 - Increased PostgreSQL Shared memory to 2048MB and worker_mem to 64MB
 - Disable playbook priority.
To disable playbook priority, edit the `/opt/cyops-workflow/sealab/sealab/config.ini` file and set the `ENABLE_PRIORITY` parameter to `false`. Then restart the `uwsgi` and `celeryd` services, using the following command:

```
# systemctl restart celeryd uwsgi
```

Notes on the tests performed

In a production environment the following factors might vary, which could affect the observations:

- The size of the alert data.
Note: For all the above tests, the average size of an alert created in FortiSOAR is 5KB.
- The number of playbooks that are being executed in parallel for each alert. For example, system playbooks for notification or triage/investigate playbooks.
- The number of steps in each playbook.
- The network bandwidth, especially for outbound connections, to applications such as VirusTotal.

Test setup for the single-node FortiSOAR appliance

This test has been invoked on a standalone FortiSOAR machine.

Tests performed

Test 1: Perform Ingestion in FortiSOAR using the FortiSIEM Ingestion Playbook

This test is executed by manually triggering FortiSIEM Ingestion playbook that creates alerts in FortiSOAR.

Steps followed

1. Created the alerts using the FortiSIEM Ingestion playbook. You can download the JSON for the sample playbooks ([Test_1_Info_Json_Files.zip](#)) so that you can run the same tests in your environment to see the performance in

your version/hardware platforms. Or, if you want to do some additions that are specific to your environment, you can also tweak the existing playbooks.

2. Once the alerts are created, measured the total time taken to create all the alerts in FortiSOAR.

Observations

The data in the following table outlines the number of alerts ingested and the total time taken to ingest those alerts.

Single Invocation Test run on a single-node FortiSOAR appliance

Number of alerts created in FortiSOAR	Total time (in seconds) taken to create all alerts in FortiSOAR	Total number of playbooks executed in FortiSOAR
1	0.32	1
5	0.47	1
10	0.67	1
25	1.69	1
50	2.66	1
100	6.13	1

Note: Once this test is completed, refer to the [pre-test conditions](#) before starting a new test.

Test 2: Perform Ingestion in FortiSOAR using the FortiSIEM Ingestion Playbook after the alerts are created and after triggering "Extraction" playbooks

This test is executed by manually triggering FortiSIEM Ingestion playbook that creates alerts in FortiSOAR. Once the alerts are created in FortiSOAR, an "Extraction" playbook is triggered and the total time taken for all the extraction playbooks to complete their execution is calculated.

Steps followed

1. Created the alerts using the FortiSIEM Ingestion playbook.
2. Once the alerts are created, the "Extraction" playbooks are triggered. You can download the JSON for the sample playbooks ([Test_2_Info_Json_Files.zip](#)) so that you can run the same tests in your environment to see the performance in your version/hardware platforms. Or, if you want to do some additions that are specific to your environment, you can also tweak the existing playbooks. The playbooks perform the following steps:
 - a. Declares variables using the "Set Variable Step".
 - b. Updates the existing indicator list using mapping.
 - c. Retrieves indicators from the source data of the alert.
 - d. Creates indicators in the "Indicators" Module.
 - e. Link alerts to the indicators.
 - f. Update Alert State.

Observations

The data in the following table outlines the number of alerts ingested, the total time taken to ingest those alerts, and the total time taken for all the triggered playbooks to complete their execution.

Single Invocation Test run on a single-node FortiSOAR appliance

Number of alerts created in FortiSOAR	Total time (in seconds) taken to create all alerts in FortiSOAR	Total number of playbooks executed in FortiSOAR
1	1.62	2
5	2.71	6
10	3.57	16
25	6.35	37
50	11.40	73
100	23.11	144

Note: Once this test is completed, refer to the [pre-test conditions](#) before starting a new test.

Test 3: Perform Ingestion in FortiSOAR using the FortiSIEM Ingestion Playbook after the alerts are created and after triggering "Extraction" and "Enrichment" playbooks

The test was executed using an automated testbed that starts FortiSIEM ingestion which in turn created alerts in FortiSOAR. Once the alerts are created in FortiSOAR, "Extraction" and "Enrichment" playbooks are triggered and the total time taken for all the extraction and enrichment playbooks to complete their execution is calculated.



The setup for this test is exactly the same, however this test additionally requires the "VirusTotal" connector to be configured.

Steps followed

1. Created the alerts using the FortiSIEM Ingestion playbook.
2. Once the alerts are created, the "Extraction" playbooks are triggered. You can download the JSON for the sample playbooks ([Test_3_Info_Json_Files.zip](#)) so that you can run the same tests in your environment to see the performance in your version/hardware platforms. Or, if you want to do some additions that are specific to your environment, you can also tweak the existing playbooks. The playbooks perform the following steps:
 - a. Declares variables using the "Set Variable Step".
 - b. Updates the existing indicator list using mapping.
 - c. Retrieves indicators from the source data of the alert.
 - d. Creates indicators in the "Indicators" Module.
 - e. Link alerts to the indicators.
 - f. Update Alert State.
3. Once the indicators were extracted, the "Enrichment" playbooks are triggered and they perform the following steps:

Once the indicators were extracted, the "Enrichment" playbooks are triggered and they perform the following steps:

- a. Matches the IP in an internal subnet through the "Utilities" subnet.
- b. Validates whether the IP is Private or Public.
- c. Performs enrichment using the "Utilities" connector, if the IP is "Private".
- d. Performs enrichment using the "VirusTotal" connector, if the IP is "Public".
- e. Updates the indicator status based on the IP's vulnerability.
- f. Updates the state of the indicator State.

Observations

The data in the following table outlines the number of alerts ingested, the total time taken to ingest those alerts, and the total time taken for all the triggered playbooks to complete their execution.

Single Invocation Test run on a single-node FortiSOAR appliance

Number of alerts created in FortiSOAR	Total time (in seconds) taken to create all alerts in FortiSOAR	Total number of playbooks executed in FortiSOAR
1	4.93	4
5	6.28	16
10	11.53	31
25	25.87	76
50	47.52	151
100	1 minute 34 seconds	301

Note: Once this test is completed, refer to the [pre-test conditions](#) before starting a new test. Also, note that enrichment of playbooks makes API calls over the Internet, and the times mentioned in this table to execute playbooks is inclusive of this time.

Sustained Invocation Test for the single-node FortiSOAR appliance

A sustenance test was also conducted with the configuration as defined in "Test 2", i.e., the test is executed by manually triggering FortiSIEM Ingestion playbook that creates alerts in FortiSOAR. Once the alerts are created in FortiSOAR, an "Extraction" playbook is triggered and the total time taken for all the extraction playbooks to complete their execution is calculated.

Number of alerts: **100/min**

Duration: **12 hours**

Playbooks configured: **As defined in Test 2 comprising of "Ingestion" and "Indicator Extraction" playbooks.**

Total number of playbooks executed: **72720**

Results


The system performed well under the sustained load. All 72000 alerts were successfully ingested and all the extraction playbooks were successfully completed without any queuing.

Graphs

The following graphs are plotted for the vital statistics for the system that was under test during the period of the test run.

CPU Load Average Utilization Graph


Analysis of CPU load average utilization when the test run was in progress on the appliance:



Using the system resources specified in the "Environment" and "Pre-Test Conditions" sections, it was observed that while the "Sustenance Test" was running the CPU utilization was normal and the performance of the system did not get impacted.

Memory Utilization Graph


Analysis of memory utilization when the test run was in progress on the appliance:



Using the system resources specified in the "Environment" and "Pre-Test Conditions" sections, it was observed that while the "Sustenance Test" was running the Memory utilization was around 10GB.

Redis PB Queue Count Graph


Analysis of Redis PB Queue Count when the test run was in progress on the appliance:



Using the system resources specified in the "Environment" and tunables configured as mentioned in the "Pre-Test Conditions" sections, it was observed that while the "Sustenance Test" the redis_pb_queue was at 0. This means that no playbooks were getting queued and all the required playbooks associated with the alerts were getting completed, i.e., alerts were created and its indicators were extracted and enriched in a minute.

IO Wait Graph


Analysis of IO Wait when the test run was in progress on the appliance:



Using the system resources specified in the "Environment" and "Pre-Test Conditions" sections, it was observed that while the "Sustenance Test" was running the IO Wait time was normal, with the average IO Wait being around 1% of the CPU idle time.

Read/Write IO Wait Graph for ElasticSearch

Analysis of Read/Write IO Wait for ElasticSearch when the test run was in progress on the appliance:




Using the system resources specified in the "Environment" and "Pre-Test Conditions" sections, it was observed that while the "Sustenance Test" was running the "Read" Wait for the ElasticSearch disk was almost 0 milliseconds. The

"Write" Wait for the ElasticSearch disk averaged around 30 milliseconds, with the maximum wait of 40 milliseconds and the minimum wait of 1 millisecond.

Read/Write IO Wait Graph for PostgreSQL

Analysis of Read/Write IO Wait for PostgreSQL when the test run was in progress on the appliance:



Using the system resources specified in the "Environment" and "Pre-Test Conditions" sections, it was observed that while the "Sustenance Test" was running, the "Read" Wait for the PostgreSQL disk averaged around 1 millisecond, with the maximum wait of 3 milliseconds and the minimum wait of 0 milliseconds. The "Write" Wait for the PostgreSQL disk averaged around 1 millisecond, with the maximum wait of 5.5 milliseconds and the minimum wait of 1 millisecond.

Single Invocation Test for the High Availability (HA) active-active cluster of two FortiSOAR nodes

Test setup for the HA active-active cluster of two FortiSOAR nodes

This test has been invoked the following setup:

- Cluster of two FortiSOAR machines that are joined in the Active-Active state using the FortiSOAR HA feature.
- The machines that form the HA cluster must be in the same network subnet.

Tests performed

Test 1: Perform Ingestion in FortiSOAR using the FortiSIEM Ingestion Playbook

This test is executed by manually triggering FortiSIEM Ingestion playbook that creates alerts in FortiSOAR.

Steps followed

1. Created the alerts using the FortiSIEM Ingestion playbook. You can download the JSON for the sample playbooks ([Test_1_Info_Json_Files.zip](#)) so that you can run the same tests in your environment to see the performance in your version/hardware platforms. Or, if you want to do some additions that are specific to your environment, you

can also tweak the existing playbooks.

2. Once the alerts are created, measured the total time taken to create all the alerts in FortiSOAR.

Observations

The data in the following table outlines the number of alerts ingested and the total time taken to ingest those alerts.

Single Invocation Test run on a two-node active-active FortiSOAR appliance

Number of alerts created in FortiSOAR	Total time (in seconds) taken to create all alerts in FortiSOAR	Total number of playbooks executed in FortiSOAR
1	0.80	1
5	0.48	1
10	0.70	1
25	1.45	1
50	2.69	1
100	6.36	1

Note: Once this test is completed, refer to the [pre-test conditions](#) before starting a new test.

Test 2: Perform Ingestion in FortiSOAR using the FortiSIEM Ingestion Playbook after the alerts are created and after triggering "Extraction" playbooks

This test is executed by manually triggering FortiSIEM Ingestion playbook that creates alerts in FortiSOAR. Once the alerts are created in FortiSOAR, an "Extraction" playbook is triggered and the total time taken for all the extraction playbooks to complete their execution is calculated.

Steps followed

1. Created the alerts using the FortiSIEM Ingestion playbook.
2. Once the alerts are created, the "Extraction" playbooks are triggered. You can download the JSON for the sample playbooks ([Test_2_Info_Json_Files.zip](#)) so that you can run the same tests in your environment to see the performance in your version/hardware platforms. Or, if you want to do some additions that are specific to your environment, you can also tweak the existing playbooks. The playbooks perform the following steps:
 - a. Declares variables using the "Set Variable Step".
 - b. Updates the existing indicator list using mapping.
 - c. Retrieves indicators from the source data of the alert.
 - d. Creates indicators in the "Indicators" Module.
 - e. Link alerts to the indicators.
 - f. Update Alert State.

Observations

The data in the following table outlines the number of alerts ingested, the total time taken to ingest those alerts, and the total time taken for all the triggered playbooks to complete their execution.

Single Invocation Test run on a two-node active-active FortiSOAR appliance

Number of alerts created in FortiSOAR	Total time (in seconds) taken to create all alerts in FortiSOAR	Total number of playbooks executed in FortiSOAR
1	3.41	2
5	2.38	6
10	3.23	16
25	5.79	37
50	10.80	73
100	20.70	144

Note: Once this test is completed, refer to the [pre-test conditions](#) before starting a new test.

Test 3: Perform Ingestion in FortiSOAR using the FortiSIEM Ingestion Playbook after the alerts are created and after triggering "Extraction" and "Enrichment" playbooks

The test was executed using an automated testbed that starts FortiSIEM ingestion which in turn created alerts in FortiSOAR. Once the alerts are created in FortiSOAR, "Extraction" and "Enrichment" playbooks are triggered and the total time taken for all the extraction and enrichment playbooks to complete their execution is calculated.



The setup for this test is exactly the same, however this test additionally requires the "VirusTotal" connector to be configured.

Steps followed

1. Created the alerts using the FortiSIEM Ingestion playbook.
2. Once the alerts are created, the "Extraction" playbooks are triggered. You can download the JSON for the sample playbooks ([Test_3_Info_Json_Files.zip](#)) so that you can run the same tests in your environment to see the performance in your version/hardware platforms. Or, if you want to do some additions that are specific to your environment, you can also tweak the existing playbooks. The playbooks perform the following steps:
 - a. Declares variables using the "Set Variable Step".
 - b. Updates the existing indicator list using mapping.
 - c. Retrieves indicators from the source data of the alert.
 - d. Creates indicators in the "Indicators" Module.
 - e. Link alerts to the indicators.
 - f. Update Alert State.
3. Once the indicators were extracted, the "Enrichment" playbooks are triggered and they perform the following steps:

Once the indicators were extracted, the "Enrichment" playbooks are triggered and they perform the following steps:

- a. Matches the IP in an internal subnet through the "Utilities" subnet.
- b. Validates whether the IP is Private or Public.
- c. Performs enrichment using the "Utilities" connector, if the IP is "Private".
- d. Performs enrichment using the "VirusTotal" connector, if the IP is "Public".
- e. Updates the indicator status based on the IP's vulnerability.
- f. Updates the state of the indicator State.

Observations

The data in the following table outlines the number of alerts ingested, the total time taken to ingest those alerts, and the total time taken for all the triggered playbooks to complete their execution.

Single Invocation Test run on a two-node active-active FortiSOAR appliance

Number of alerts created in FortiSOAR	Total time (in seconds) taken to create all alerts in FortiSOAR	Total number of playbooks executed in FortiSOAR
1	5.39	4
5	6.46	16
10	8.59	31
25	19.44	76
50	35.66	151
100	1 minute 6 seconds	301

Note: Once this test is completed, refer to the [pre-test conditions](#) before starting a new test. Also, note that enrichment of playbooks makes API calls over the Internet, and the times mentioned in this table to execute playbooks is inclusive of this time.

Sustained Invocation Test for the HA active-active cluster of two FortiSOAR appliance

A sustenance test was also conducted with the configuration as defined in "Test 2", i.e. the test is executed by manually triggering FortiSIEM Ingestion playbook that creates alerts in FortiSOAR. Once the alerts are created in FortiSOAR, an "Extraction" playbook is triggered and the total time taken for all the extraction playbooks to complete their execution is calculated.

Number of alerts: **100/min**

Duration: **12 hours**

Playbooks configured: **As defined in Test 2 comprising of "Ingestion" and "Indicator Extraction" playbooks.**

Total number of playbooks executed: **72720**

Results

The system performed well under the sustained load. All 72000 alerts were successfully ingested and all the extraction playbooks were successfully completed without any queuing.

Graphs


The following graphs are plotted for the vital statistics for the HA cluster that was under test during the period of the test run.



All the graphs included in this section are from the Primary/Active Node.

CPU Load Average Utilization Graph


Analysis of CPU load average utilization when the test run was in progress on the appliance:



Using the system resources specified in the "Environment" and "Pre-Test Conditions" sections, it was observed that while the "Sustenance Test" was running the CPU utilization was normal and the performance of the system did not get impacted.

Memory Utilization Graph


Analysis of memory utilization when the test run was in progress on the appliance:



Using the system resources specified in the "Environment" and "Pre-Test Conditions" sections, it was observed that while the "Sustenance Test" was running the Memory utilization was around 10GB.

Redis PB Queue Count Graph


Analysis of Redis PB Queue Count when the test run was in progress on the appliance:



Using the system resources specified in the "Environment" and tunables configured as mentioned in the "Pre-Test Conditions" sections, it was observed that while the "Sustenance Test" the redis_pb_queue was at 0. It went to 4 only once and again came back to 0. This means that no playbooks were getting queued and all the required playbooks associated with the alerts were getting completed in a minute.

IO Wait Graph


Analysis of IO Wait when the test run was in progress on the appliance:



Using the system resources specified in the "Environment" and "Pre-Test Conditions" sections, it was observed that while the "Sustenance Test" was running the IO Wait time was normal, with the average IO Wait being around 1% of the CPU idle time.

Read/Write IO Wait Graph for ElasticSearch


Analysis of Read/Write IO Wait for ElasticSearch when the test run was in progress on the appliance:



Using the system resources specified in the "Environment" and "Pre-Test Conditions" sections, it was observed that while the "Sustenance Test" was running the "Read" Wait for the ElasticSearch disk was almost 0 milliseconds. The "Write" Wait for the ElasticSearch disk averaged around 30 milliseconds, with the maximum wait of 40 milliseconds and the minimum wait of 1 millisecond.

Read/Write IO Wait Graph for PostgreSQL

Analysis of Read/Write IO Wait for PostgreSQL when the test run was in progress on the appliance:



Using the system resources specified in the "Environment" and "Pre-Test Conditions" sections, it was observed that while the "Sustenance Test" was running, the "Read" Wait for the PostgreSQL disk averaged around 0 millisecond. The "Write" Wait for the PostgreSQL disk averaged around 8 milliseconds, with the maximum wait of 25 milliseconds and the minimum wait of 1 millisecond.



Copyright© 2021 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, FortiCare® and FortiGuard®, and certain other marks are registered trademarks of Fortinet, Inc., in the U.S. and other jurisdictions, and other Fortinet names herein may also be registered and/or common law trademarks of Fortinet. All other product or company names may be trademarks of their respective owners. Performance and other metrics contained herein were attained in internal lab tests under ideal conditions, and actual performance and other results may vary. Network variables, different network environments and other conditions may affect performance results. Nothing herein represents any binding commitment by Fortinet, and Fortinet disclaims all warranties, whether express or implied, except to the extent Fortinet enters a binding written contract, signed by Fortinet's General Counsel, with a purchaser that expressly warrants that the identified product will perform according to certain expressly-identified performance metrics and, in such event, only the specific performance metrics expressly identified in such binding written contract shall be binding on Fortinet. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. In no event does Fortinet make any commitment related to future deliverables, features or development, and circumstances may change such that any forward-looking statements herein are not accurate. Fortinet disclaims in full any covenants, representations, and guarantees pursuant hereto, whether express or implied. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable.