

FortiSIEM - Sizing Guide

FORTINET DOCUMENT LIBRARY

<http://docs.fortinet.com>

FORTINET VIDEO GUIDE

<http://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

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

FEEDBACK

Email: techdocs@fortinet.com



TABLE OF CONTENTS

Change Log	4
FortiSIEM Sizing Information	5
Minimum Requirements.....	6
Internal Scalability Tests.....	7
Test Setup.....	7
Test Success Criteria.....	7
Hardware Appliance EPS Test.....	8
Virtual Appliance EPS Test with FortiSIEM Event Database.....	9
Virtual Appliance EPS Test with Elasticsearch Database.....	10
Recommended Sizing for FortiSIEM Event DB Based Deployment.....	12
Processing Requirement.....	12
Storage Requirement for FortiSIEM EventDB.....	13
Recommended Sizing for Elasticsearch Based Deployment.....	15
Processing Requirement.....	15
Storage Requirement for Elasticsearch.....	16

Change Log

Date	Change Description
03/30/2018	Initial version of FortiSIEM Sizing Guide.
04/12/2018	Revision 2 with updates to Storage Requirements for FortiSIEM EventDB and Elasticsearch Data Nodes sections.

FortiSIEM Sizing Information

This document provides information about the following:

[Minimum Requirements](#)

[Internal Scalability Tests](#)

[Hardware Appliance EPS Test](#)

[Virtual Appliance EPS Test with FortiSIEM Event Database](#)

[Virtual Appliance EPS Test with Elasticsearch Database](#)

[Recommended Sizing for FortiSIEM Event DB Based Deployment](#)

[Recommended Sizing for Elasticsearch Based Deployment](#)

Minimum Requirements

Browser Display

FortiSIEM, like most monitoring, SIEM and analytics tools, shows a lot of information on the screen at once. FortiSIEM HTML GUI has chosen a bigger font for legibility reasons. Hence, we recommend that users have a minimum 1680x1050 desktop display resolution.

Hardware

Minimum hardware requirements for FortiSIEM nodes are as follows.

Node	CPU	RAM	Local Disk
Supervisor	8 vCPU	24 GB (32 GB is using Elasticsearch)	200 GB (80 GB for OS and App, 60 GB for CMDB and 60 GB for SVN)
Worker	4 vCPU	16 GB	200 GB (80 GB for OS and App; rest for used)
Collector	2 vCPU	4 GB	40 GB

- Supervisor VA needs more memory since it hosts many heavy-duty components such as Application Server (Java), PostgreSQL Database Server and Rule Master.
- With Elasticsearch, Supervisor VA also hosts the Java Query Server component for communicating with Elasticsearch – hence the need for additional 8 GB memory.

Note that these are only the minimum requirements. The performance may improve by increasing vCPUs and RAM in certain situations. External storage depends on your EPS mix and the number of days of log storage needs. To provide more meaningful guidance, scalability tests were conducted as described below.

Internal Scalability Tests

FortiSIEM team performed several scalability tests described below.

Test Setup

- A specific set of events were sent repeatedly to achieve the target EPS.
- The target EPS was constant over time.
- A set of Linux servers were monitored via SNMP and performance monitoring data was collected.
- Events triggered many incidents.

Test Success Criteria

The following success criteria should be met on testing:

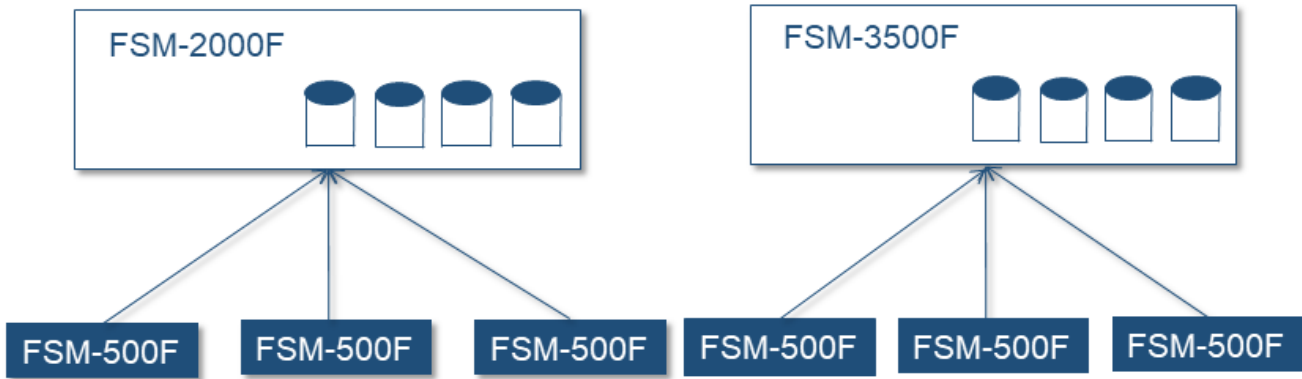
- Incoming EPS must be sustained without any event loss.
- Summary dashboards should be up to date and not fall behind.
- Widget dashboards should show data indicating that inline reporting is keeping up.
- Incidents should be up to date.
- Real-time search should show current data and trend chart should reflect incoming EPS.
- GUI navigation should be smooth.
- CPU, memory and IOPS are not maxed out. Load average must be less than the number of cores.

The tests were run for three cases:

- All-in-one FSM Hardware Appliance: FSM-2000F and FSM-3500F with collectors FSM-500F sending events.
- FSM Virtual Appliance with FortiSIEM EventDB as the data store.
- FSM Virtual Appliance with Elasticsearch as the data store.

Hardware Appliance EPS Test

The test beds were as follows:



The results are shown below:

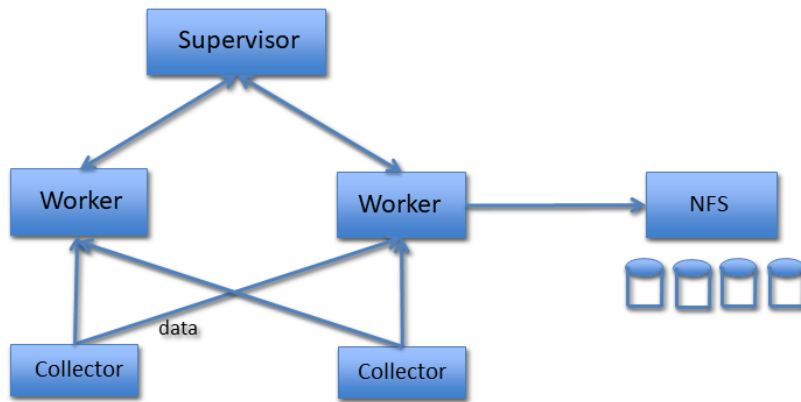
FortiSIEM HW Appliance	Event Sender			Sustained EPS without Loss
	Collector Model	Count	EPS/Collector	
FSM-2000F	FSM-500F	3	5K	15K
FSM-3500F	FSM-500F	4	8K	30K

Virtual Appliance EPS Test with FortiSIEM Event Database

All tests were done in AWS. The following hardware was used.

Type	AWS Instance Type	Hardware Spec
Collector	c4.xlarge	4vCPU, 7 GB RAM
Worker	c4.2xlarge	8vCPU, 15 GB RAM
Super	m4.4xlarge	16vCPU, 64 GB RAM, CMDB Disk 10K IOPS
NFS Server	c4.2xlarge	8vCPU, 16 GB RAM, 10K IOPS

The test bed is as follows:



The following result shows 10K EPS sustained per Worker with over 20K CMDB Devices.

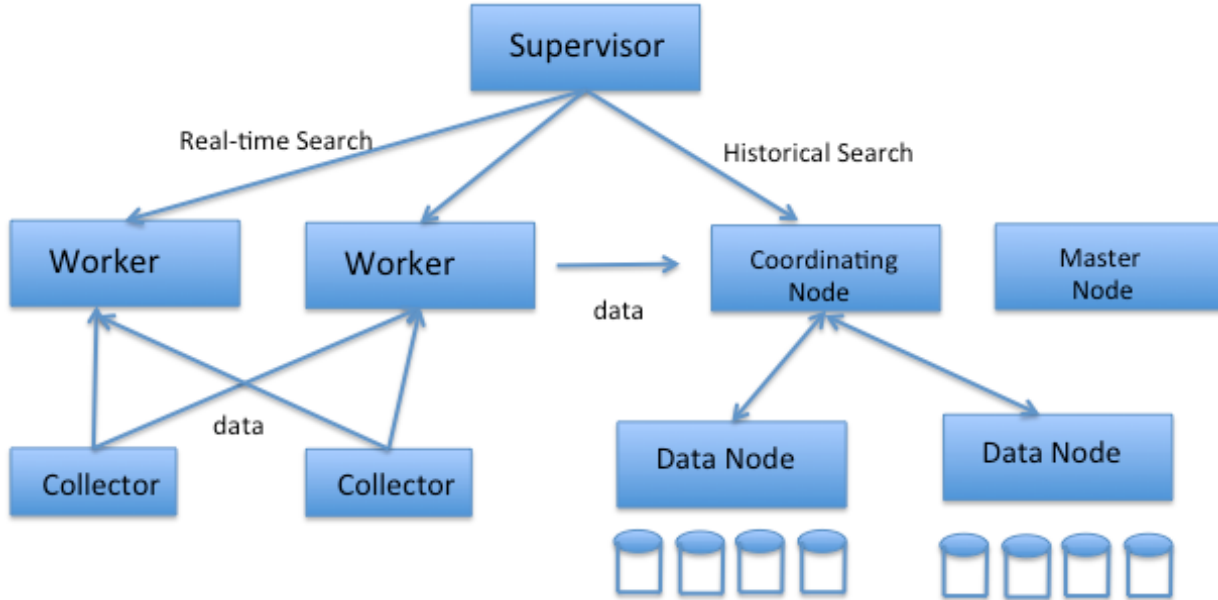
Event Sender			Event Handler				
Collector Count	EPS/Collector	Monitored Device/Collector	Super	Workers	Orgs	CMDB Device	Sustained EPS without Loss
150	200	150	1	3	150	22,500	30K

Virtual Appliance EPS Test with Elasticsearch Database

All tests were done in AWS. The following hardware was used.

Type	AWS Instance Type	Hardware Spec
Collector	c4.xlarge	4vCPU, 7 GB RAM
Worker	c4.2xlarge	8vCPU, 15 GB RAM
Super	m4.4xlarge	16vCPU, 64 GB RAM, CMDB Disk 10K IOPS
Elastic Search Master Node	c3.2xlarge	8vCPU, 16 GB RAM with 8 GB JVM
Elastic Search Coordinating Node	m5.4xlarge	16vCPU, 64 GB RAM with 30 GB JVM allocation
Elastic Search Data Node	i3.4xlarge	16vCPU, 122 GB RAM, 1.9TBx2 NVMe SSD Instance-store Volumes, 30 GB JVM

The test bed was as follows:



The following result shows 5K EPS sustained per Data Node with over 20K CMDB Devices.

Event Sender			Event Handler					
Collector Count	EPS/Collector	Monitored Device/Collector	Super	Workers	Elastic (M/CO/DN/Shards)*	Orgs	CMDB Device	Sustained EPS without Loss
150	200	150	1	3	1/1/5/10	150	22,500	30K

* M = Elasticsearch Master, CO = Elasticsearch Co-ordinator, DN = Elasticsearch Data Node

Recommended Sizing for FortiSIEM Event DB Based Deployment

Processing Requirement

Requirement		Recommendation			
EPS	Deployment	HW Model	SW Configuration		
			Nodes	HW Per Node (vCPU, RAM)	NFS IOPS
Up to 5K	Hardware	FSM-2000F			
Up to 5K	Software		All-in-one	16,24GB	
5K – 10K	Hardware	FSM-2000F			
5K – 10K	Software		Supervisor	16,24GB	
			1 Worker	8,16GB	2000
10K – 15K	Hardware	FSM-3500F			
10K – 15K	Software		Supervisor	16,24GB	
			2 Workers	8,16GB	3000
15K – 25K	Hardware	FSM-3500F			
15K – 25K	Software		Supervisor	16,24GB	
			3 Workers	16,16GB	5000
25K – 35K	Software		Supervisor	16,24GB	
			4 Workers	16,16GB	7000
Add 10K EPS	Software		Add 1 Worker	16,16GB	Add 2000 IOPS

Storage Requirement for FortiSIEM EventDB

FortiSIEM storage requirement depends on three factors:

- EPS
- Bytes/log mix in your environment
- Compression ratio (8:1)

You are likely licensed for Peak EPS. Typically, EPS peaks during morning hours on weekdays and goes down dramatically after 2 pm on weekdays and also remains low on weekends. So the average EPS should be used to calculate storage needs.

For calculating Bytes/log, consider the following aspects:

- Network devices and Linux servers tend to send shorter logs (150-200 bytes/log) while Windows Security logs tend to be much larger (500-1000 bytes/log).
- Busy corporate firewalls and domain controllers tend to send much higher log volumes (higher EPS) than other systems, assuming they are sending all logs.
- Database indices built on logs for efficient searching consumes significant storage as well.
- ASCII text (syslog) compresses much better than binary (for example, Netflow)

Therefore, it is difficult to properly assume a specific Bytes/log mix in your environment without measurement. Our experience from sampling of 5 large customers has shown that Bytes/log is between 100-150 including all factors – device mix, log mix, indexing cost and compression. We calculated this by dividing the total FortiSIEM event file size (in `\data`) over one day by the total number of events on that day, and then averaging over a few days.

The table below shows two scenarios – Worst case and Average case for NFS storage. In Worst case, Peak EPS and 150 Bytes/log is used. In the Average case, 0.5 Peak EPS and 100 Bytes/log is used.

Peak EPS	Storage (Months)	NFS Storage (TB)*	
		Worst case	Average case
1000	12	5	1.66
1000	24	9	3
1000	36	14	4.66
2000	12	9	3
2000	24	19	6.33
2000	36	28	9.33
5000	12	23	7.66

Peak EPS	Storage (Months)	NFS Storage (TB)*	
		Worst case	Average case
5000	24	47	15.66
5000	36	70	23.33
10000	12	47	15.66
10000	24	93	31
10000	36	140	46.66

NFS Storage (TB):

- Worst case = $(\text{Peak EPS} * 150 * 86400 * 30 * \text{Storage(Months)}) / 10^{12}$
- Average case = $(0.5 * \text{Peak EPS} * 100 * 86400 * 30 * \text{Storage(Months)}) / 10^{12}$

Recommended Sizing for Elasticsearch Based Deployment



Adding or moving shards is easy but splitting is not possible. Plan ahead for shard sizing is very important.

Processing Requirement

Requirement	Recommendation				
EPS	ES Configuration	Hardware per node (vCPU, RAM)	Elastic JVM RAM	Shards	Replica
Up to 1K - without Replica	All-in-one	(8,16GB)	8GB	5	0
Up to 1K - with Replica	3 node cluster	(8,16GB)	8GB	5	1
1K-5K - with Replica	3 node cluster	(8,64GB)	30GB	5	1
5K-10K - with Replica	Coordinating and Master Node	(8,32GB)	16GB		
	3 Data Nodes	(8,64GB)	30GB	5	1
10K-15K - with Replica	Coordinating Node	(16,32GB)	16GB		
	Master Node	(8,16GB)	8GB		
	3 Data Nodes	(16,64GB)	30GB	10	1
15K-25K - with Replica	Coordinating Node	(16,64GB)	30GB		
	Master Node	(8,16GB)	8GB		
	5 Data Nodes	(16,64GB)	30GB	15	1
25K-35K - with Replica	Coordinating Node	(16,64GB)	30GB		

Requirement	Recommendation				
EPS	ES Configuration	Hardware per node (vCPU, RAM)	Elastic JVM RAM	Shards	Replica
	Master Node	(8,16GB)	8GB		
	7 Data Nodes	(16,64GB)	30GB	20	1
35K-45K - with Replica	Coordinating Node	(16,64GB)	30GB		
	Master Node	(8,16GB)	8GB		
	9 Data Nodes	(16,64GB)	30GB	25	1
Add 5K EPS - with Replica	Add 1 Data Node	(16,64GB)	30GB	Add 3 Shards	1

Storage Requirement for Elasticsearch

Elasticsearch consumes more storage than NFS because it indexes the data more heavily than FortiSIEM event database.

FortiSIEM Elasticsearch storage requirement depends on two factors:

- EPS
- Bytes/log mix in your environment

You are likely licensed for Peak EPS. Typically, EPS peaks during morning hours on weekdays and goes down dramatically after 2 pm on weekdays and also remains low on weekends. So the average EPS should be used to calculate storage needs.

For calculating Bytes/log, consider the following aspects:

- Network devices and Linux servers tend to send shorter logs (150-200 bytes/log) while Windows Security logs tend to be much larger (500-1000 bytes/log).
- Busy corporate firewalls and domain controllers tend to send much higher log volumes (higher EPS) than other systems, assuming they are sending all logs.
- Database indices built on logs for efficient searching consumes significant storage as well.
- ASCII text (syslog) compresses much better than binary (for example, Netflow)

Therefore, it is difficult to properly assume a specific Bytes/log mix in your environment without measurement. Our internal scalability test environment shows Bytes/log is around 1000 including all factors – device mix, log mix, indexing cost and compression. We calculated this by dividing the total Elasticsearch database file size (in \data) over one day by the total number of events on that day, and then averaging over a few days.

The table below shows two scenarios – Worst case and Average case for Storage/Cluster. In Worst case, Peak EPS and 1000 Bytes/log is used. In the average case, 0.5 Peak EPS is used. As we gather experience with more customers, we will publish average Bytes/log and update the Average storage requirements.

Peak EPS	Replica	Storage (Months)	Storage/Cluster (TB)	
			Worst case	Average case
1000	0	12	31	15.5
1000	1	12	62	31
2000	1	12	124	62
5000	1	12	311	155.5
10000	1	6	311	155.5
15000	1	6	467	233.5
25000	1	3	389	194.5
50000	1	3	778	389

Storage per Cluster (TB):

- Worst case = $(\text{Peak EPS} * 1000 * 86400 * \text{Storage(Month)} * 30 * (\text{Replica} + 1)) / 10^{12}$
- Average case = $(0.5 * \text{Peak EPS} * 1000 * 86400 * \text{Storage(Month)} * 30 * (\text{Replica} + 1)) / 10^{12}$

