

# Sizing Guide - Elasticsearch

FortiSIEM 7.5.1



**FORTINET DOCUMENT LIBRARY**

<https://docs.fortinet.com>

**FORTINET VIDEO LIBRARY**

<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/training-certification>

**FORTINET TRAINING INSTITUTE**

<https://training.fortinet.com>

**FORTIGUARD LABS**

<https://www.fortiguard.com>

**END USER LICENSE AGREEMENT**

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

**FEEDBACK**

Email: [techdoc@fortinet.com](mailto:techdoc@fortinet.com)



05/04/2026

FortiSIEM 7.5.1 Sizing Guide - Elasticsearch

# TABLE OF CONTENTS

<b>Change Log</b> .....	<b>4</b>
<b>FortiSIEM Sizing Guide - Elasticsearch</b> .....	<b>5</b>
Minimum Requirements .....	5
Hardware .....	5
Internal Scalability Tests .....	6
Test Setup .....	6
Test Success Criteria .....	6
Virtual Appliance EPS Test with Elasticsearch Database .....	7
Sizing Online Deployment .....	8
Background .....	9
Recommended Elasticsearch Configuration .....	10
Sizing of Coordinator Only Nodes .....	10
Sizing of Hot Data Nodes .....	11
Sizing of Warm Data Nodes .....	12
Sizing of Frozen Data Nodes .....	13
Sizing Archive Deployments .....	14
EventDB Based Deployments .....	14
HDFS Based Deployments .....	15
References .....	17

# Change Log

Date	Change Description
12/23/2025	Sizing Guide - Elasticsearch release for 7.5.0.
05/04/2026	Sizing Guide - Elasticsearch release for 7.5.1.

# FortiSIEM Sizing Guide - Elasticsearch

This document provides information about the following topics:

- [Minimum Requirements](#)
  - [Hardware](#)
- [Internal Scalability Tests](#)
  - [Test Setup](#)
  - [Test Success Criteria](#)
  - [Virtual Appliance EPS Test with Elasticsearch Database](#)
- [Sizing Online Deployment](#)
  - [Background](#)
  - [Recommended Elasticsearch Configuration](#)
  - [Sizing of Coordinator Only Nodes](#)
  - [Sizing of Hot Data Nodes](#)
  - [Sizing of Warm Data Nodes](#)
  - [Sizing of Frozen Data Nodes](#)
- [Sizing Archive Deployments](#)
  - [EventDB Based Deployments](#)
  - [HDFS Based Deployments](#)
- [References](#)

## Minimum Requirements

### Hardware

Minimum hardware requirements for FortiSIEM nodes are as follows.

Node	vCPU	RAM	Local Disks
Supervisor (All in one)	Minimum – 12 Recommended - 32	Minimum <ul style="list-style-type: none"> <li>• without UEBA – 24GB</li> <li>• with UEBA - 32GB</li> </ul> Recommended <ul style="list-style-type: none"> <li>• without UEBA – 32GB</li> <li>• with UEBA - 64GB</li> </ul>	OS – 25GB OPT – 100GB CMDB – 60GB SVN – 60GB Local Event database – based on need
Supervisor (Cluster)	Minimum – 12 Recommended - 32	Minimum <ul style="list-style-type: none"> <li>• without UEBA – 24GB</li> </ul>	OS – 25GB OPT – 100GB

Node	vCPU	RAM	Local Disks
		<ul style="list-style-type: none"> <li>with UEBA - 32GB</li> </ul> Recommended	CMDB – 60GB SVN – 60GB
Workers	Minimum – 8 Recommended - 16	Minimum – 16GB Recommended <ul style="list-style-type: none"> <li>without UEBA – 24GB</li> <li>with UEBA - 32GB</li> </ul>	OS – 25GB OPT – 100GB
Collector	Minimum – 4 Recommended – 8 (based on load)	Minimum – 4GB Recommended – 8GB	OS – 25GB OPT – 100GB

- 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.
- For OPT - 100GB, the 100GB disk for /opt will consist of a single disk that will split into 2 partitions, /OPT and swap. The partitions will be created and managed by FortiSIEM when configFSM.sh runs.

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 the following cases:

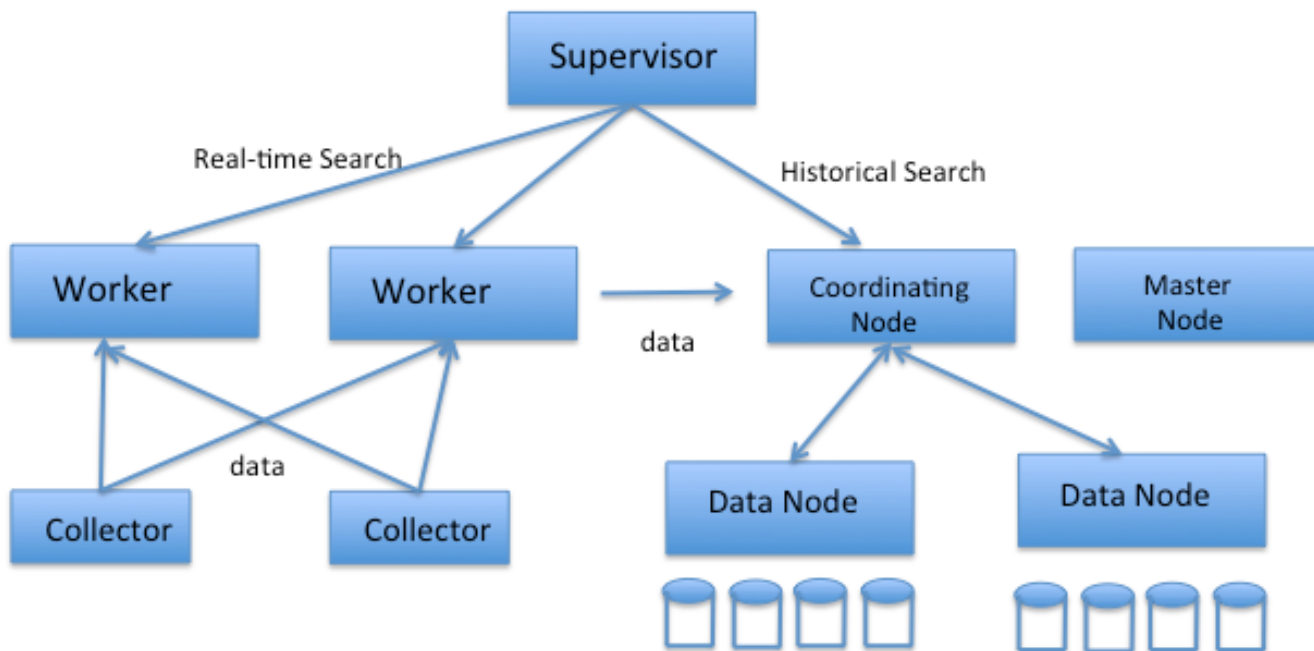
- All-in-one FSM Hardware Appliance: FSM-2000F and FSM-3500F with collectors FSM-500F sending events.
- FSM Virtual Appliance with Elasticsearch as the data store.

## 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, CMDDB 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

## Sizing Online Deployment

This section provides information about the following Elasticsearch based deployment topics:

- Background
- Recommended Elasticsearch Configuration
- Sizing of Coordinator Only Nodes
- Sizing of Hot Data Nodes
- Sizing of Warm Data Nodes
- Sizing of Frozen Data Nodes

## Background

An Elasticsearch deployment consists of

- Master node (required)
- Coordinator Only nodes (required)
- Data nodes – Hot, Warm and Frozen (See below)

Keep the following points in mind about Hot, Warm and Frozen Data nodes:

- FortiSIEM inserts events into Hot nodes, so Hot nodes need fast disk I/O to handle inserts and reads. A solid state drive (SSD) with more than 200Gb/s I/O throughput are recommended for hot nodes.
- User can specify Hot node retention policy (days) in FortiSIEM. When this limit is reached, or Hot node disk usage reaches its high watermark (75% full), events are moved from Hot nodes to Warm nodes, to make room for new events in the Hot node. Warm nodes only handle event reads, so they can afford slightly lower speed disks than Hot nodes. Hard disk drives or DAS/SAN disks with about 100 Gb/s I/O throughput are recommended. Warm nodes will have similar query response times as Hot nodes, as they aren't performing expensive indexing operations.
- User can specify Warm node retention policy (days) in FortiSIEM. When this limit is reached, or Warm node disk usage reaches its high watermark (75% full), events are moved from Warm nodes to Frozen nodes. In the Frozen nodes, indices are flushed from memory to disk, so Frozen nodes can have much larger disks. Fortinet recommends Frozen nodes to have similar disk speeds as Warm nodes, but with larger capacity. When a user queries for data in Frozen nodes, the Frozen nodes will temporarily rehydrate the required indices to memory, and then flush to disk after the query is complete. Therefore, Frozen nodes will have higher query response time.

When using Elasticsearch in FortiSIEM, hot nodes are the minimum requirement. A budget friendly way to add additional storage is to add Warm nodes. For even more capacity, Frozen nodes can be added.

While Elasticsearch Frozen nodes can provide some Archive options, FortiSIEM provides two additional Archive options with higher compression – FortiSIEM EventDB on NFS and HDFS (see the Sizing Guide for EventDB for more information). Should you choose to use archiving, ensure you use the **Real Time Archive** option so that events are written to the Archive at the point of insertion into Hot nodes. This eliminates the need for moving data from Online to Archive, as it is prohibitively expensive to read events out of Elasticsearch and write back to Archive.

To summarize, the following options are available for Elasticsearch:

- Online
  - Elasticsearch
    - Hot node only
    - Hot node and Warm node
    - Hot node, Warm node, and Frozen node
- Archive
  - EventDB on NFS
  - HDFS

FortiSIEM has a dynamic shard management feature to keep Elasticsearch working optimally.

- For Hot nodes, FortiSIEM lets a shard grow up to 40GB (Elasticsearch limit 50GB), and closes it once the 40GB limit is reached. In a low EPS situation, a shard can span multiple days. FortiSIEM uses an aliasing technique to string the shards together. In a high EPS situation, there can be multiple shards in a day.

- Segment merge is performed to reduce heap usage for indices older than 2 days.

## Recommended Elasticsearch Configuration

Fortinet recommends the following configuration steps.

- Replica – at least 1
- Master, Coordinator Only nodes and Data nodes on different machines
- 3 Master nodes – each with 8 vCPU, 16 GB RAM
- At least 2 Coordinator Only nodes – each with 16 vCPU, 32 GB RAM. Two Coordinator Only nodes are chosen for failover. The exact number of Coordinator Only nodes depends on the EPS. See below for details.
- At least 3 Hot Data nodes – each with 32 vCPU, 64GB RAM and SSD disks with at least 200 Gb/s I/O throughput. The exact number of Hot Data nodes depends on the EPS and retention policy (see below).
- If you decide to deploy Warm nodes, deploy at least 3 Warm Data nodes – each with 32 vCPU, 64GB RAM and disks with at least 100 Gb/s I/O throughput. The exact number of Warm Data nodes depends on retention policy (see below).
- If you decide to deploy Frozen nodes, deploy at least 3 Frozen Data nodes – each with 16 vCPU, 64GB RAM and around 100 Gb/s I/O throughput. The exact number of Frozen Data nodes depends on retention policy (see below).
- If you decide to utilize the Archive option, then choose the **Real Time Archive** option.
- Enable FortiSIEM dynamic shard management feature (Under **Shard Allocation**, select **Dynamic**).
- Do not choose **Per Org Index** unless you have to. Keeping a separate index per organization enables you to efficiently delete an organization’s data, but increases the shard count, which in turn requires more nodes because of the 500 shards per node limit.
- Follow the [Pre-install considerations in the setup guide](#).

## Sizing of Coordinator Only Nodes

Our testing has shown that 1 Coordinator Only node with 16 vCPU and 64 GB RAM (32 GB to Elasticsearch and the rest to the operating system) can handle 200K EPS. So choose the number of Coordinator Only nodes based on your EPS and keep one extra in case a Coordinator Only node dies.

Peak EPS	Coordinator Only Nodes (16 vCPU and 64 GB RAM)
50K	2
100K	2
200K	3
500K	4
1 Million	6

## Sizing of Hot Data Nodes

Choose each Hot Data node to have 32 vCPU, 64GB RAM (32 GB to Elasticsearch and the rest to the operating system), and SSD with at least 200Gbps I/O throughput. Use the calculations below to find the exact number of Hot nodes. This depends on the EPS and the retention policy.

### Assumptions

- Our experiments have shown that FortiSIEM can insert 60K EPS per Data node without Replica and without any loss
- Measurements from various SIEM installs have shown that Elasticsearch consumes an average of 500 bytes to store an event and all its parsed attributes
- Memory to Disk Ratio = 1:30. See the Elasticsearch reference documents in [References](#).
- Max 75% disk capacity for storing event indices since Elasticsearch will use the rest of the 25%. See the Elasticsearch reference documents in [References](#).

### Storage per Day

Suppose

- R: Number of Replica (at least 1 is recommended)
- E: EPS
- D: Retention (days) in Hot nodes

Data per day =  $E * \text{\#seconds in a day (86400)} * 500 \text{ Bytes} * (R + 1)$

Storage per day =  $1.25 * \text{Data per day}$

### # Hot Data Nodes from Elasticsearch Constraint

Min # Hot Data nodes =  $\text{Storage per day} * D / \text{RAM} / 30$

### # Hot Data Nodes from EPS Constraint

Suppose

- R: Number of Replica (at least 1 is recommended)
- E: EPS

Since FortiSIEM can insert 60K EPS without Replica and without any loss,

Min # Hot Data nodes =  $E * (R+1) / 60K$

### Cluster Wide Shard Count Limit

In Elasticsearch 6.x, Fortinet has observed that Elasticsearch CLI performance degrades when the total number of shards in the cluster (including Hot and Warm nodes) is more than 15K. Newer versions may have a higher upper limit.

You can calculate the number of shards as follows.

Node Count: N

Disk Size on each node: D TB

Assuming 40GB per shard, the total number of shards =  $(N * D \text{ TB}) / 40\text{GB}$

For example, at 50K EPS: (consult the two tables below)

Hot Node: 7 days retention, number of shards =  $(19 * 2\text{TB}) / 40\text{GB} = 950$

Warm Node: 30 days retention, number of shards =  $(15 * 10\text{TB}) / 40\text{GB} = 3750$

Total shards in the cluster = 4700

As you map out higher EPS and longer days, make sure the number of shards stays within the limit.

### Examples

Here are some representative numbers for Replica = 1

EPS	Storage per Day	Retention (Days)	Hot Data Nodes (32vCPU, 64GB RAM, SSD)	
			Node Count	Disk Size
10K	1TB	7	4	2TB
		14	8	2TB
		30	16	2TB
50K	5TB	7	19	2TB
		14	38	2TB
		30	80	2TB
100K	10TB	7	38	2TB
		14	76	2TB
		30	160	2TB
200K	20TB	7	76	2TB
		14	152	2TB
		30	320	2TB
500K	50TB	7	190	2TB
		14	380	2TB
		30	760	2TB
1 Million	100TB	7	380	2TB
		14	760	2TB
		30	1520	2TB

## Sizing of Warm Data Nodes

Warm nodes are configured identically as Hot Data nodes, except that memory to disk ratio = 1:160. Make sure you have 32vCPU, 64GB RAM to host with 32 GB to Elasticsearch and 32GB to the operating system. See the Elasticsearch reference documents in [References](#).

Min # Warm Data nodes = Storage per day \* D / RAM / 160

## Examples

Here are some representative numbers for Replica = 1

EPS	Storage per Day	Retention (Days)	Warm Data Nodes (32vCPU, 64GB RAM and ~100Gbps Disk I/O)	
			Node Count	Disk Size
10K	1TB	30	3	10TB
		60	6	10TB
		90	9	10TB
50K	5TB	30	15	10TB
		60	30	10TB
		90	45	10TB
100K	10TB	30	30	10TB
		60	60	10TB
200K	20TB	30	60	10TB
		60	120	10TB
500K	50TB	30	150	10TB
		60	300	10TB
1 Million	100TB	30	300	10TB
		60	600	10TB

## Sizing of Frozen Data Nodes

Frozen nodes are configured identically as Warm or Hot Data nodes, except for a higher memory to disk ratio = 1:1000. Note that the required query latency, rather than system resources, often limits the amount of data stored on Frozen nodes. See the Elasticsearch reference documents in [References](#).

Frozen nodes behave identically except for memory to disk ratio = 1:1000

Min # Frozen Data nodes = Storage per day \* D / RAM / 1000

## Examples

Here are some representative numbers for Replica = 1

EPS	Storage per Day	Retention (Days)	Frozen Data Nodes (16vCPU, 64GB RAM, HDD <100Gbps)	
			Node Count	Disk Size
10K	1TB	90	2	60TB
		180	4	60TB
		365	7	60TB
50K	5TB	90	8	60TB
		180	16	60TB
		365	32	60TB
100K	10TB	90	15	60TB
		180	30	60TB
		365	60	60TB
200K	20TB	90	30	60TB
		180	60	60TB
		365	120	60TB
500K	50TB	90	75	60TB
		180	150	60TB
		365	300	60TB
1 Million	100TB	90	150	60TB
		180	300	60TB
		365	600	60TB

## Sizing Archive Deployments

For Elasticsearch, events can be archived to NFS using EventDB format or to HDFS.

- [EventDB Based Deployments](#)
- [HDFS Based Deployments](#)

## EventDB Based Deployments

In this situation, online workers are used to query the Archived EventDB database, so only a NFS infrastructure is required. Since Archived data is not indexed, our experiments have shown that Archived EventDB needs

about 60% storage compared to Online EventDB. This information can be used to estimate the amount of NFS storage required for Archive.

EPS	Retention	NFS Storage	
		Worst Case (100 Bytes/log)	Average Case (66 Bytes/log)
5K	6 months	7.5 TB	2.5 TB
	1 year	15 TB	5 TB
	3 years	45 TB	15 TB
10K	6 months	15 TB	5 TB
	1 year	30 TB	10 TB
	3 years	90 TB	30 TB
20K	6 months	30 TB	10 TB
	1 year	60 TB	20 TB
	3 years	180 TB	60 TB
50K	6 months	75 TB	25 TB
	1 year	150 TB	50 TB
	3 years	450 TB	150 TB
100K	6 months	150 TB	50 TB
	1 year	300 TB	100 TB
	3 years	900 TB	300 TB

Worst Case Storage = EPS \* 86400 \* worst case bytes/log \* retention

Average Case Storage = 0.5 \* EPS \* 86400 \* average case bytes/log \* retention

Used 1024 for B -> KB etc.

## HDFS Based Deployments

An HDFS based deployment needs the following:

- Name node
- Spark Master node
- Spark Slave node
- Data node

According to HDFS best practices, the following co-locations are possible:

- Node Type A containing Name node – need 2 of these – 8vCPU and 16GB RAM each
- Node Type B containing Spark Master and Slave node – need 2 of these -- 8vCPU and 16GB RAM each

- Node Type C containing Slave node and Data node – need  $N$  of these based on insert and query requirements.

Our experiments have shown that

- HDFS Parquet file system uses 125 bytes/event (25% of Elasticsearch storage and 30% more than FortiSIEM EventDB).
- HDFS can insert events at 200K EPS per Data node, so insert speed is not a determining criteria.

## Spark / HDFS Resource Allocation Considerations

Spark / HDFS can run the following 4 jobs.

1. Real time Archive job
2. Non-real time Archive job
3. Query Job from Supervisor
4. Merge job to merge small files written to HDFS into bigger files

The resource allocation strategy is as follows:

$totalCore$  = Total number of cores on Spark Worker nodes

$totalWorker$  = Total number of Spark Worker nodes

$availCore$  = ( $totalCore - totalWorker$ ) available to the 4 job types above

$nodeCore$  = Total number of cores on one single Spark worker node

### Strategy

1. Real time archive job is allocated  $rm\_lambda\_percent$  (default 30) of resources. This parameter is defined in `/home/admin/FortiSIEM/config.txt`.
2. Remaining resources are allocated to 3 other job types in an equal manner. The parameter  $rm\_max\_concurrent\_jobs$  (default 5) defined in `/home/admin/FortiSIEM/config.txt` defines the maximum number of such jobs. Each executor needs to be allocated 5 cores to run, Hence the maximum value of  $rm\_max\_concurrent\_jobs$  should be the minimum value of the following two calculations:
  - a.  $(availCore * ((100 - rm\_lambda\_percent) / 100)) / 5$ .
  - b.  $(nodeCore - 1) * ((100 - rm\_lambda\_percent) / 100) / 5 * totalWorker$

These two calculations may differ as the lower integer needs to be retrieved when dividing by 5.

### Example

10 Spark nodes each 16 core and 32GB RAM.

$totalCore$  = 160

$totalWorker$  = 10

$availCore$  = 150

$rm\_lambda\_percent$  = 30

$rm\_max\_concurrent\_job$  = 20

1. Real time archive job will be assigned with at most 9 executors, each with 5 core and 10GB RAM. Spark will dynamically adjust the executor numbers based on running status.
2. A resource pool of 20 jobs are created, each with 5 core and 10GB RAM. Each of these can run the other 3 jobs based on demand.

Based on this, the following sizing is suggested.

EPS	Storage per Day (Replica = 1)	Retention	Total Storage	Spark Slave + Data Node (16 vCPU, 32GB RAM, HDD – 100 Gbps)	
				Count	Disk Storage
10K	125 GB	1 year	45 TB	6	8 TB
		3 years	134 TB	6	24 TB
50K	625 GB	1 year	225 TB	10	24 TB
		3 years	675 TB	10	80 TB
100K	1.25 TB	1 year	460 TB	10	48 TB
		3 years	1,380 TB	10	160 TB
200K	2.5 TB	1 year	920 TB	10	96 TB
		3 years	2,760 TB	10	300 TB
500K	5.75 TB	1 year	2,250 TB	8	300 TB
		3 years	6,750 TB	25	300 TB
1 Million	11.5 TB	1 year	4,450 TB	15	300 TB
		3 years	13,350 TB	25	600 TB

## References

Elasticsearch Concepts and Sizing Guide

<https://www.elastic.co/pdf/elasticsearch-sizing-and-capacity-planning.pdf>

Elasticsearch Sizing Guide

<https://www.elastic.co/blog/benchmarking-and-sizing-your-elasticsearch-cluster-for-logs-and-metrics>



[www.fortinet.com](http://www.fortinet.com)

Copyright© 2026 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, FortiCare® and FortiGuard®, and certain other marks are registered trademarks of Fortinet, Inc., 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 Chief Legal Officer, 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. 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.