



Sizing Guide - ClickHouse

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



05/04/2026

FortiSIEM 7.5.1 Sizing Guide - ClickHouse

TABLE OF CONTENTS

Change Log	4
FortiSIEM Sizing Guide - ClickHouse	5
Minimum Requirements	5
Hardware	5
Internal Scalability Tests	6
Test Setup	6
Test Success Criteria	6
Hardware Appliance EPS Test with ClickHouse	7
Virtual Appliance EPS Test with ClickHouse Database	8
Sizing Online Deployment	9
Processing Requirement	9
Storage Requirement	15
Configuring ClickHouse/Migrating Event Database to ClickHouse	17

Change Log

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

FortiSIEM Sizing Guide - ClickHouse

This document provides information about the following topics:

- [Minimum Requirements](#)
 - [Hardware](#)
- [Internal Scalability Tests](#)
 - [Test Setup](#)
 - [Test Success Criteria](#)
 - [Hardware Appliance EPS Test with ClickHouse](#)
 - [Virtual Appliance EPS Test with ClickHouse Database](#)
- [Sizing Online Deployment](#)
 - [Processing Requirement](#)
 - [Storage Requirement](#)
- [Configuring ClickHouse/Migrating Event Database to ClickHouse](#)

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"> • without UEBA – 24GB • with UEBA - 32GB Recommended <ul style="list-style-type: none"> • without UEBA – 32GB • with UEBA - 64GB 	OS – 25GB OPT – 100GB CMDB – 60GB SVN – 60GB ClickHouse DB - based on EPS and retention
Supervisor (Cluster)	Minimum – 12 Recommended - 32	Minimum <ul style="list-style-type: none"> • without UEBA – 24GB • with UEBA - 32GB Recommended <ul style="list-style-type: none"> • without UEBA – 32GB • with UEBA - 64GB 	OS – 25GB OPT – 100GB CMDB – 60GB SVN – 60GB ClickHouse DB - based on EPS and retention
Workers (Data Node)	Minimum – 16 Recommended - 32	Minimum – 32GB Recommended	OS – 25GB OPT – 100GB

Node	vCPU	RAM	Local Disks
		<ul style="list-style-type: none"> without UEBA – 64GB with UEBA - 64GB 	ClickHouse DB - based on EPS and retention
Workers (Keeper Only Node)	Minimum 8 Recommended 16	Minimum - 16GB Recommended 16 GB	OS – 25GB OPT – 100GB Data - 200GB
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.
- 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.

Hardware Appliance EPS Test with ClickHouse

The test bed is shown below. Scripts generated events on FSM-500F Collectors, which parsed those events and sent to the appliances.

Appliance	Hardware Spec	Event Sender			Sustained EPS without Loss
		Collector Model	Count	EPS/Collector	
FSM-2000F	<ul style="list-style-type: none"> • 12vCPU (1x6C2T) • 32GB RAM • 12x3TB SATA (3 RAID Groups) 	FSM-500F	3	5K	15K
FSM-2000G	<ul style="list-style-type: none"> • 40vCPU (2x10C2T) • 128GB RAM • 4x1TB SSD (RAID5) • 8x4TB SAS (2 RAID50 Groups) 	FSM-500F	6	7K	20K
FSM-2200G	<ul style="list-style-type: none"> • 40vCPU (2x10C20T) • 128GB RAM • 4x1.92TB SSD (RAID5) • 8x4TB SAS (2 RAID50 Groups) 	FSM-500F	6	7K	40K
FSM-3500G	<ul style="list-style-type: none"> • 48vCPU (2x12C2T) • 128GB RAM • 24x4TB SATA (3 RAID50 Groups) 	FSM-500F	6	8K	40K
FSM-3600G	<ul style="list-style-type: none"> • 48vCPU (2x16C32T) • 128GB RAM • 4x3.84TB SATA SSD • 12x8TB SATA HDD (3 RAID50 Groups) 	FSM-500F	8	8K	60K

Notes:

1. Event Ingestion speed increased two fold in FSM-2000G with ClickHouse compared to FortiSIEM EventDB. ClickHouse event database made better utilization of the vCPUs in the system.
2. The FSM-2000F recommended sustained EPS from version 7.1.0 is 7,500 EPS. FortiSIEM 7.x releases add new capabilities, such as the Machine Learning frameworks that require additional compute resources. Operating FSM-2000F at or below the recommended sustained EPS provides spare performance capacity for day-to-day SOC activity that should be considered beyond EPS ingestion performance alone.
3. For FortiSIEM 3500G, the insert performance of FortiSIEM EventDB and ClickHouse is identical as FortiSIEM EventDB could also use disk striping for better I/O.

Virtual Appliance EPS Test with ClickHouse Database

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

Node Type	AWS Instance Type	Hardware Specification
Collector	c5.2xlarge	8 vCPU, 16 GB.
Worker as ClickHouse Keeper node	C6a.8xlarge	32 vCPU, 64 GB, SSD 125Mbps throughput
Worker as ClickHouse Data/Query Node	C6a.8xlarge	32 vCPU, 64 GB, SSD 1GBps throughput
Supervisor	m6a.8xlarge	32 vCPU, 128 GB, CMDB Disk 10K IOPS

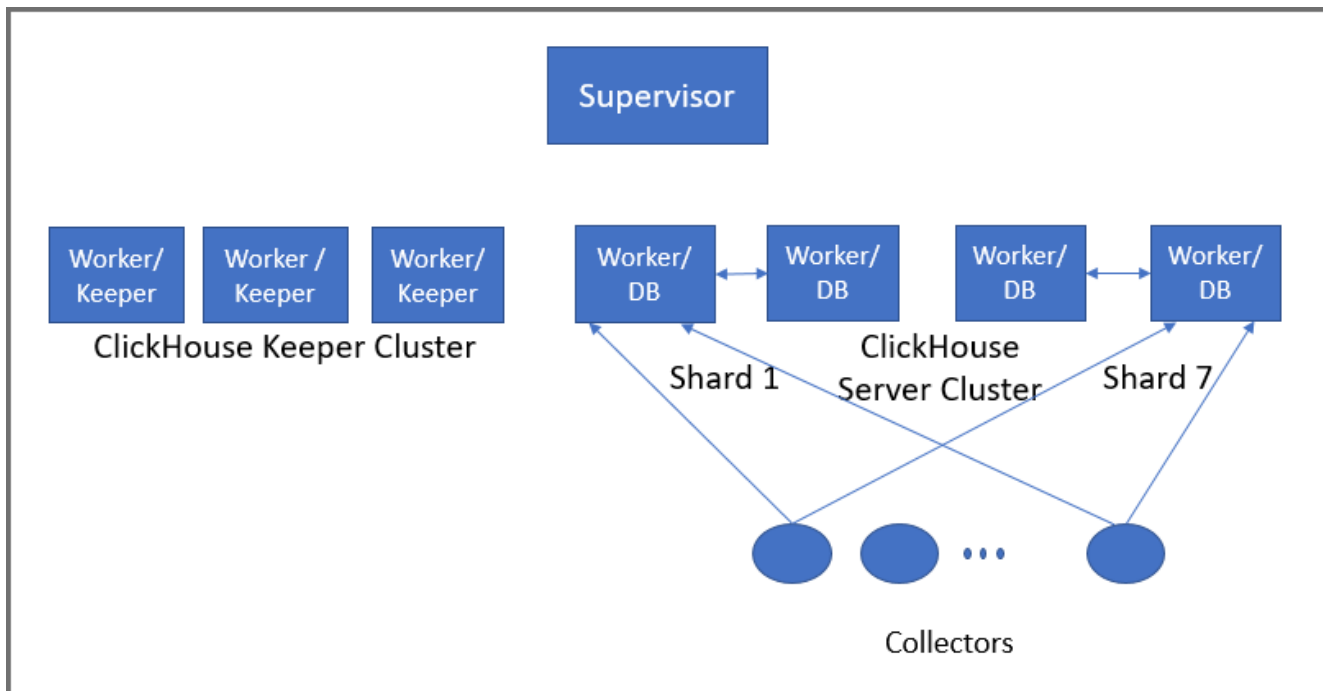
Based on the requirement to handle 500K EPS, the following setup was used:

- 1 Supervisor
- 3 Worker nodes as part of ClickHouse Keeper Cluster
- 14 Worker nodes as part of ClickHouse Server Cluster
 - 7 shards
 - 2 Workers in each shard. This means that 2 copies of each event were kept (Replication = 2).
- 150 Collectors, each sending 3.3K EPS to the 14 Workers in the ClickHouse Server Cluster, in a round robin fashion. Each Worker replicated its received events to the other Worker within the same shard.
- Collectors could also send events to the ClickHouse Keeper Cluster nodes, but this was not done. The ClickHouse Keeper Cluster nodes were dedicated to Replication management.
- Each Worker handles 35.7K EPS.

See [ClickHouse Configuration](#) in the latest Online Help for details on setting up ClickHouse Clusters.

See the testbed below. Scripts generated events on the Collectors, which were sent to the Workers. Service provider deployment was used. There were 150 Organizations and each Collector belonging to an Organization discovered and monitored the performance of 150 other Collectors in other Organization. This resulted in 22.5K devices in CMDB and each were being discovered using SNMP and monitored for basic performance metrics including CPU, Memory, Disk and Network interface utilization.

500K EPS were sustained without any event loss for over 2 days. 5 users logged on the system and ran queries and visited various parts of the user interface.



Sizing Online Deployment

Processing Requirement

- Hardware Appliance Deployments
- Software Based Deployments

Hardware Appliance Deployments

EPS	Deployment	Replication	Hardware Model	Network
0-20K	Hardware	1	2000F, 2000G, 2200G, 3500G, 3600G	1Gbps
20K-40K	Hardware	1	3500G, 3600G	1Gbps
40K-50K	Hardware	1	3600G	1Gbps

Software Based Deployments

Software based deployments can be scaled out to handle more EPS by adding shards and adding Worker nodes in each shard. See [ClickHouse Operational Overview](#) for details. Follow these principles for a stable deployment:

1. Whenever possible, deploy separate ClickHouse Keeper nodes. This is true especially at medium to high EPS or you will run into many concurrent heavy-duty queries. In these cases, Keeper functionality may compete for CPU, Memory, and Disk I/O resources with Insert and Query. If Keeper does not get resources, replication will stop, database will become read only and event insertion stops. In the table below, Fortinet recommends **3 dedicated Keeper nodes** for 60K EPS and above. For 20K-60K, dedicated Keeper nodes is an option.
2. If more than 50% Keeper nodes are lost, then RAFT protocol quorum is lost and database may become read only, and event insertion stops. For this reason, Fortinet recommends **3 Keeper nodes** whenever possible as it can sustain 1 lost node.
 - a. If you run 2 Keeper nodes, then loss of 1 node causes quorum to be lost and database may become read only.
 - b. If you run 1 Keeper node, then loss of 1 node causes complete loss of Keeper cluster and database may become read only.

In both these cases, follow the steps in [Recovering from Losing Quorum](#) to recover from lost quorum or complete keeper cluster loss. Using more than 3 Keeper nodes may lead to increased replication overhead.

3. **Use SSD for Hot Tier**, especially for medium to high EPS. This will speed up event insertion and queries.
4. If you need to handle more EPS, then add more shards, using the table below as a guide.
5. If you need to make queries run faster, there are two options:
 - a. Add more shards
or
 - b. Add more Data + Query nodes in existing shards

Both these approaches will spread out the data to more nodes.

Requirement		Configuration	
EPS	Replication	Supervisor/Worker Hardware	ClickHouse Topology
0-5K	1 (meaning 1 copy of events)	1 Supervisor – 16vCPU, 24GB RAM, 200MBps Disk	1 Shard with 1 Replica The Shard has Supervisor with Data and Query flag checked. Supervisor is also Keeper node
0-5K	2 (meaning 2 copies of events)	1 Supervisor – 16vCPU, 24GB RAM, 200MBps Disk 1 Worker – 16vCPU, 24GB RAM, 200MBps Disk 1 Gbps Network	1 Shard with 2 Replicas The Shard has Supervisor and Worker with both Data and Query flags checked. Supervisor is also Keeper Node
5K-10K	1	1 Supervisor – 32vCPU, 32GB RAM, 200MBps Disk	1 Shard with 1 Replica The Shard has Supervisor with both Data and Query flags checked. Supervisor is also Keeper node
5K-10K	2	1 Supervisor – 16vCPU, 32GB RAM, 200MBps Disk	1 Shard with 2 Replicas The Shard has Supervisor and Worker with both Data and Query flags checked.

Requirement		Configuration	
EPS	Replication	Supervisor/Worker Hardware	ClickHouse Topology
		1 Worker – 16vCPU, 32GB RAM, 200MBps Disk 1 Gbps Network	Supervisor is also Keeper Node
10K-20K	1	1 Supervisor - 48vCPU, 64GB RAM, 200MBps Disk	1 Shard with 1 Replica The Shard has Supervisor with both Data and Query flags checked. Supervisor is also Keeper node
10K-20K	2	1 Supervisor – 32vCPU, 64GB RAM, 200MBps Disk 1 Worker – 32vCPU, 64GB RAM, 200MBps Disk 1 Gbps Network	1 Shard with 2 Replicas The Shard has Supervisor and Worker with both Data and Query flags checked. Supervisor is also Keeper Node
20K-30K	1	1 Supervisor – 32vCPU, 64GB RAM, 200MBps Disk 1 Worker – 32vCPU, 64GB RAM, 200MBps Disk 1 Gbps Network	1 Shard with 1 Replica The Shard has Worker with both Data and Query flags checked. Supervisor is Keeper node
20K-30K	2	1 Supervisor – 32vCPU, 64GB RAM, 200MBps Disk 2 Workers – 32vCPU, 64GB RAM, 200MBps Disk 1 Gbps Network	1 Shard with 2 Replicas The Shard has 2 Workers with both Data and Query flags checked. Supervisor is also Keeper Node
		1 Supervisor – 32vCPU, 64GB RAM, 200MBps Disk 2 Workers – 32vCPU, 64GB RAM, 200MBps Disk 3 Workers – 16vCPU, 16GB RAM, 200MBps Disk 1 Gbps Network	1 Shard with 2 Replicas The Shard has 2 Workers with both Data and Query flags checked. 3 Workers (16vCPU) acting as Keeper only
30K-60K	2	1 Supervisor – 32vCPU, 64GB RAM, 200MBps Disk 2 Workers – 32vCPU, 64GB RAM, 500MBps Disk 1 Worker – 16vCPU, 16GB RAM, 200MBps Disk 10Gbps Network	1 Shard with 2 Replicas Each shard – 2 (32vCPU) Workers with both Data and Query flags checked. 1 Worker (16vCPU) acting as Keeper only

Requirement		Configuration	
EPS	Replication	Supervisor/Worker Hardware	ClickHouse Topology
		1 Supervisor – 32vCPU, 64GB RAM, 200MBps Disk 2 Workers – 32vCPU, 64GB RAM, 500MBps Disk 3 Workers – 16vCPU, 16GB RAM, 200MBps Disk 10Gbps Network	1 Shard with 2 Replicas Each shard – 2 (32vCPU) Workers with both Data and Query flags checked. 3 Workers (16vCPU) acting as Keeper only
60K-125K	2	1 Supervisor – 32vCPU, 64GB RAM, 200MBps Disk 4 Workers – 32vCPU, 64GB RAM, 500MBps Disk 3 Workers – 16vCPU, 16GB RAM, 200MBps Disk 10Gbps Network	2 Shards with 2 Replicas per shard Each shard has 2 (32vCPU) Workers with both Data and Query flags checked. 3 (16vCPU) Workers acting as dedicated Keeper Nodes
125K-175K	2	1 Supervisor – 32vCPU, 64GB RAM, 200MBps Disk 6 Workers – 32vCPU, 64GB RAM, 500MBps Disk 3 Workers – 16vCPU, 16GB RAM, 200MBps Disk 10Gbps Network	3 Shards with 2 Replicas per shard Each shard has 2 (32vCPU) Workers with both Data and Query flags checked. 3 (16vCPU) Workers acting as dedicated Keeper Nodes
175K-250K	2	1 Supervisor – 32vCPU, 64GB RAM, 200MBps Disk 8 Workers – 32vCPU, 64GB RAM, 500MBps Disk 3 Workers – 16vCPU, 16GB RAM, 200MBps Disk 10Gbps Network	4 Shards with 2 Replicas per shard Each shard has 2 (32vCPU) Workers with both Data and Query flags checked. 3 (16vCPU) Workers acting as dedicated Keeper Nodes
250K-300K	2	1 Supervisor – 32vCPU, 64GB RAM, 200MBps Disk 10 Workers – 32vCPU, 64GB RAM, 500MBps Disk 3 Workers – 16vCPU, 16GB RAM, 200MBps Disk 10Gbps Network	5 Shards with 2 Replicas per shard Each shard has 2 (32vCPU) Workers with both Data and Query flags checked. 3 (16vCPU) Workers acting as dedicated Keeper Nodes
300K-360K	2	1 Supervisor – 32vCPU, 64GB RAM, 200MBps Disk	6 Shards with 2 Replicas per shard

Requirement		Configuration	
EPS	Replication	Supervisor/Worker Hardware	ClickHouse Topology
		12 Workers – 32vCPU, 64GB RAM, 500MBps Disk 3 Workers – 16vCPU, 16GB RAM, 200MBps Disk 10Gbps Network	Each shard has 2 (32vCPU) Workers with both Data and Query flags checked. 3 (16vCPU) Workers acting as dedicated Keeper Nodes
360K-420K	2	1 Supervisor – 32vCPU, 64GB RAM, 200MBps Disk 14 Workers – 32vCPU, 64GB RAM, 1GBps Disk 3 Workers – 16vCPU, 16GB RAM, 200MBps Disk 10Gbps Network	7 Shards with 2 Replicas per shard Each shard has 2 (32vCPU) Workers with both Data and Query flags checked. 3 (16vCPU) Workers acting as dedicated Keeper Nodes
420K-500K	2	1 Supervisor – 32vCPU, 64GB RAM, 200MBps Disk 16 Workers – 32vCPU, 64GB RAM, 1GBps Disk 3 Workers – 16vCPU, 16GB RAM, 200MBps Disk 10Gbps Network	8 Shards with 2 Replicas per shard Each shard has 2 (32vCPU) Workers with both Data and Query flags checked. 3 (16vCPU) Workers acting as dedicated Keeper Nodes
500K-550K	2	1 Supervisor – 32vCPU, 64GB RAM, 500MBps Disk 18 Workers – 32vCPU, 64GB RAM, 1GBps Disk 3 Workers – 16vCPU, 16GB RAM, 200MBps Disk 10Gbps Network	9 Shards with 2 Replicas per shard Each shard has 2 (32vCPU) Workers with both Data and Query flags checked. 3 (16vCPU) Workers acting as dedicated Keeper Nodes
550K-600K	2	1 Supervisor – 32vCPU, 64GB RAM, 500MBps Disk 20 Workers – 32vCPU, 64GB RAM, 1GBps Disk 3 Workers – 16vCPU, 16GB RAM, 200MBps Disk 10Gbps Network	10 Shards with 2 Replicas per shard Each shard has 2 (32vCPU) Workers with both Data and Query flags checked. 3 (16vCPU) Workers acting as dedicated Keeper Nodes
600K-650K	2	1 Supervisor – 32vCPU, 64GB RAM, 500MBps Disk 22 Workers – 32vCPU, 64GB RAM, 1GBps Disk	11 Shards with 2 Replicas per shard Each shard has 2 (32vCPU) Workers with both Data and Query flags checked.

Requirement		Configuration	
EPS	Replication	Supervisor/Worker Hardware	ClickHouse Topology
		3 Workers – 16vCPU, 16GB RAM, 200MBps Disk 10Gbps Network	3 (16vCPU) Workers acting as dedicated Keeper Nodes
650K-700K	2	1 Supervisor – 32vCPU, 64GB RAM, 500MBps Disk 24 Workers – 32vCPU, 64GB RAM, 1GBps Disk 3 Workers – 16vCPU, 16GB RAM, 200MBps Disk 10Gbps Network	12 Shards with 2 Replicas per shard Each shard has 2 (32vCPU) Workers with both Data and Query flags checked. 3 (16vCPU) Workers acting as dedicated Keeper Nodes
700K-750K	2	1 Supervisor – 32vCPU, 64GB RAM, 500MBps Disk 26 Workers – 32vCPU, 64GB RAM, 1GBps Disk 3 Workers – 16vCPU, 16GB RAM, 200MBps Disk 10Gbps Network	13 Shards with 2 Replicas per shard Each shard has 2 (32vCPU) Workers with both Data and Query flags checked. 3 (16vCPU) Workers acting as dedicated Keeper Nodes
750K-800K	2	1 Supervisor – 32vCPU, 64GB RAM, 500MBps Disk 28 Workers – 32vCPU, 64GB RAM, 1GBps Disk 3 Workers – 16vCPU, 16GB RAM, 200MBps Disk 10Gbps Network	14 Shards with 2 Replicas per shard Each shard has 2 (32vCPU) Workers with both Data and Query flags checked. 3 (16vCPU) Workers acting as dedicated Keeper Nodes
800K-850K	2	1 Supervisor – 32vCPU, 64GB RAM, 500MBps Disk 30 Workers – 32vCPU, 64GB RAM, 1GBps Disk 3 Workers – 16vCPU, 16GB RAM, 200MBps Disk 10Gbps Network	15 Shards with 2 Replicas per shard Each shard has 2 (32vCPU) Workers with both Data and Query flags checked. 3 (16vCPU) Workers acting as dedicated Keeper Nodes
850K-900K	2	1 Supervisor – 32vCPU, 64GB RAM, 500MBps Disk 32 Workers – 32vCPU, 64GB RAM, 1GBps Disk 3 Workers – 16vCPU, 16GB RAM, 200MBps Disk	16 Shards with 2 Replicas per shard Each shard has 2 (32vCPU) Workers with both Data and Query flags checked. 3 (16vCPU) Workers acting as dedicated Keeper Nodes

Requirement		Configuration	
EPS	Replication	Supervisor/Worker Hardware	ClickHouse Topology
		10Gbps Network	
900K-950K	2	1 Supervisor – 32vCPU, 64GB RAM, 500MBps Disk 34 Workers – 32vCPU, 64GB RAM, 1GBps Disk 3 Workers – 16vCPU, 16GB RAM, 200MBps Disk 10Gbps Network	17 Shards with 2 Replicas per shard Each shard has 2 (32vCPU) Workers with both Data and Query flags checked. 3 (16vCPU) Workers acting as dedicated Keeper Nodes
950K-1M	2	1 Supervisor – 32vCPU, 64GB RAM, 500MBps Disk 36 Workers – 32vCPU, 64GB RAM, 1GBps Disk 3 Workers – 16vCPU, 16GB RAM, 200MBps Disk 10Gbps Network	18 Shards with 2 Replicas per shard Each shard has 2 (32vCPU) Workers with both Data and Query flags checked. 3 (16vCPU) Workers acting as dedicated Keeper Nodes

For more than 1 million EPS, contact FortiSIEM Professional Services.

VM Collector Performance

Collector EPS performance will vary based on the overall load applied to the Collector, which may include event pulling jobs and performance monitoring. Typically, a Collector that is focused on events can sustain 8K EPS with 8 vCPU, 8GB Memory.

Storage Requirement

FortiSIEM event storage requirement depends on the following factors:

- Events per second (EPS)
- Bytes/event
- Compression Ratio
- Retention Period

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.

Bytes/event depends on the rate of event types found in your environment. Unix and Router logs tend to be in the 200-300 Bytes range, Firewall logs (e.g. Fortinet, Palo Alto) tend to be in the 700-1,500 Bytes range, Windows Security logs tend to be a little larger (1,500 – 2,000 Bytes), and Cloud logs tend to be much larger (2,000 Bytes -10K Bytes sometimes).

Fortinet has chosen Zstandard (ZSTD) compression algorithm for ClickHouse event database. The overall compression ratio depends on:

- Size of raw events
- Number of attributes parsed from a raw event. Parsed attributes add storage overhead, but they are needed for searches to work efficiently. Parsing a raw event during search would slow down searches considerable. FortiSIEM also adds about 20-30 meta data fields such as geo-location including country, city, longitude, latitude for source/destination/reporting IP fields, when such fields are found in events.
- Number of string valued attributes in the raw event. String valued attributes typically provide better compression.

It is best for the user to estimate or measure the EPS and Bytes/event for their environment. If you have stored a sufficient mix of events in a file, then you can count Bytes/event as the file size divided by the number of lines in that file.

The compression provided by FortiSIEM varies with event size and number of parsed and stored fields. Compression is higher for larger events of 1,000 Bytes or more and lower for smaller events. For example, a compression ratio of 15:1 is generally seen for logs over 1000 bytes and 25 parsed fields.

The storage requirement can be calculated as follows: EPS * Bytes/event * Compression ratio * Retention period (remember to normalize the units).

Example 1:

The following example illustrates a general storage requirement.

- Suppose in your environment that the peak EPS is 10K, and average EPS is 2K. An estimated EPS may be 6K.
- Average Raw Bytes/event is 500 Bytes
- Compression ratio 10:1
- Retention period 2 weeks (14 days) in Hot storage and 2.5 months (76 days) in Warm storage
- Replication = 2 (meaning 2 copies of data)

Then

- Storage per day: $(2 * 6000 * 86400 * 500) / (10 * 1024 * 1024 * 1024)$ GB = 48.3GB. The general formula is:
Storage per day = $(\text{Replication} * \text{EPS} * \text{Seconds in a day} * (\text{Bytes/Event})) / (\text{Compression} * 1024 * 1024 * 1024)$ GB
- Hot storage requirement for 14 days
 - Cluster wide: 676GB
 - Assuming 1 shard and 2 Data/Query Nodes per shard, per node storage is 338GB
- Warm storage requirement for 76 days
 - Cluster Wide: 3.58TB
 - Assuming 1 shard and 2 Data/Query Nodes per shard, per node storage is 1.79TB

Example 2:

This example illustrates the storage requirements for a larger deployment.

- Suppose in your environment that the peak EPS is 100K, and average EPS is 50K. An estimated EPS may be 75K.

- Average Raw Bytes/event is 1200 Bytes
- Compression ratio 15:1
- Retention period 30 days in Hot storage and 365 days in Warm storage
- Replication = 2 (meaning 2 copies of data)

Then

- Storage per day: $(2 * 75000 * 86400 * 1200) / (15 * 1024 * 1024 * 1024)$ GB = 965.6GB. The general formula is: Storage per day = $(\text{Replication} * \text{EPS} * \text{Seconds in a day} * (\text{Bytes/Event})) / (\text{Compression} * 1024 * 1024 * 1024)$ GB
- Hot storage requirement for 30 days
 - Cluster wide: 28.29TB
 - Assuming 2 shards and 2 Data/Query Nodes per shard, per node storage is 7.08TB
- Warm storage requirement for 365 days
 - Cluster Wide: 344.18TB
 - Assuming 2 shards and 2 Data/Query Nodes per shard, per node storage is 86.05TB

Configuring ClickHouse/Migrating Event Database to ClickHouse

If you would like to configure ClickHouse for FortiSIEM, see [Configuring ClickHouse Based Deployments](#) for more information.

If you have an existing EventDB and would like to migrate to ClickHouse, see [EventDB to ClickHouse](#) for more information.



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.