



Administration Guide

FortiAnalyzer-BigData 7.4.1



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



August 12, 2025

FortiAnalyzer-BigData 7.4.1 Administration Guide

58-741-1024788-20250812

TABLE OF CONTENTS

About FortiAnalyzer-BigData	6
Main Features	6
Key terms and concepts	7
FortiAnalyzer-BigData hardware environment	11
Connect to the FortiAnalyzer-BigData CLI	12
GUI overview	13
Cluster Manager	13
Custom refresh settings	14
Commands management	14
Notifications management	16
Host management	17
CLI Remote Console	18
Role assignment	19
Service management	21
Service groups	22
Service details	23
Monitor	26
Dashboards	26
Filtering the Dashboard	27
Customizing the Dashboard	27
Logs and metrics	29
Explore logs	30
Explore metrics	34
Health	36
Health Check	36
Alert	37
Hyperscale firewall logging support	45
Set up Security Manager hosts external IP addresses	45
Configure FortiAnalyzer-BigData as log server on hyperscale FortiGate	46
Configure FortiAnalyzer-BigData Hyperscale Ingestion	48
Device Manager and log rate	48
Search Hyperscale log in Log View	50
Global search	51
Starting a global search	51
Global Search settings	51
Log types (Global Search)	52
Histogram	53
Event Inspector	53
Faceted Search (+, -, focus)	53
Time Window	55
Search History	55
Split View	56

Live Streaming Search	57
Cross-Cluster Search Federation	57
Create a new Search Federation (Example)	58
Search with Federation	63
Log Query Language (LogQL)	64
Log Stream Selector	64
Log Pipeline	65
Job management and automation	67
Job history	68
Built-in automation jobs	69
Custom automation jobs	69
Custom job templates	71
Data management	75
Manage storage policy	78
Data backup	79
Incremental backups	82
Data restore	84
Migrate an ADOM to another Storage Pool	86
Truncate historical logs of the migrated ADOM	90
Bootloader	91
Bootloader Main Page	92
1. Configure Network	92
2. Install OS	93
3. Set Role	93
4. Set Chassis ID	94
5. Set Blade ID	94
6. Reset OS	94
7. Reset OS and Clear User Data	95
8. Upgrade Bootloader	95
9. Extract hardware detailed information by lshw	95
10. Check and repair hard drives	95
11. Reset root password	95
12. Reboot	95
sh. shell	96
General maintenance and best practices	97
Backup and restore to external HDFS	97
Schedule maintenance tasks for off-peak hours	97
Graceful system shutdown	97
Maintain database integrity	98
Upgrade FortiAnalyzer-BigData	99
Scaling FortiAnalyzer-BigData	103
How to scale out	103
Set up Main Host HA in stacked FortiAnalyzer-BigData 4500F chassis	104
How to remove a chassis from a stacked setup	107
Remove an extender chassis	107

Reset FortiAnalyzer-BigData	109
Soft reset FortiAnalyzer-BigData	109
Hard reset FortiAnalyzer-BigData	110
Data-at-rest encryption on FortiAnalyzer-BigData 4500F	111
Initialize data-at-rest encryption	111
Open encrypted data disk partitions	112
After a blade power cycle	112
After a graceful chassis power cycle	113
When performing an upgrade or soft reset	114
When scaling out the cluster or replacing a blade	115
Troubleshooting	117
Safe Mode	117
What to do if an upgrade fails	118
What to do if a soft reset fails	118
What to do if a hard reset fails	119
How to repair disk failures	119
How to replace a blade on a FortiAnalyzer-BigData appliance	120
How to reset a single host	121
How to rebalance the data	121
How to recover from an unhealthy service status	122
Core services	122
Data Lake services	123
Message Broker services	124
How to recover from a full disk	125
How to fix Kudu metadata corruption	125
How to enable/disable PXE boot server on Security Event Manager Controller	127
How to collect Diagnostic Logs	128
FAQ	129
Change Log	132

About FortiAnalyzer-BigData

FortiAnalyzer-BigData improves upon base FortiAnalyzer appliances and offers analytics-powered security and event log management to process large volumes of data. FortiAnalyzer-BigData is redesigned with a new distributed backend and high-end hardware. The Security Event Manager, the backend log engine of FortiAnalyzer-BigData, is a horizontally scalable, high availability (HA) system that supports the needs of large enterprise organizations. The Security Event Manager comprises multiple server blades working together as a cluster, so you can add new blades to expand and scale the Security Event Manager as your organization grows.

Main Features

FortiAnalyzer-BigData offers the following features:

High ingestion throughput

A single FortiAnalyzer-BigData can sustain 300k events per seconds (EPS) log ingestion.

FortiAnalyzer-BigData can sustain high throughput ingestion while continuing to perform analytics workload in the background.

Horizontal scalability backend

You can add additional appliance chassis to a running FortiAnalyzer-BigData without shutting down the system. This allows you to scale out the storage and query throughput.

Built-in high availability and fault tolerant backend

The backend, Security Event Manager, offers out-of-box fault tolerance and high availability with no need for initial configuration. All running services run under an active HA mode where data is replicated three times into different data hosts.

Easily recoverable data

By following regular backup scheduling procedures, you can recover lost data. FortiAnalyzer-BigData's backup drive configuration works with external Hadoop Distributed File System (HDFS) URLs.

Ease of management

FortiAnalyzer-BigData has a new Cluster Manager tile so you can manage and set up FortiAnalyzer-BigData from a centralized location. You can also monitor various service metrics, current host status, server logs and more from the Cluster Manager GUI.

Key terms and concepts

This section contains key terms used in FortiAnalyzer-BigData.

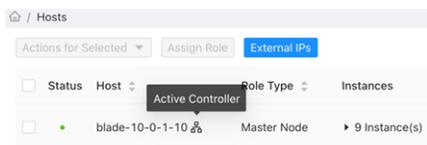
Security Event Manager

The Security Event Manager is the cluster formed by multiple server blades that serve the web GUI and performs the workload for data processing, persistence, query, and management of security log events.

Security Event Manager Controller

The Security Event Manager Controller, or cluster controller, is a single host within the Security Event Manager that functions as the main controller for the hosts. This host is responsible for the DHCP, configuration management, and lifecycle management such as upgrades, resets, and more. If this host goes down, it can automatically failover to a standby host.

To find out which of the hosts is the active Controller host, go to the Host view in the Cluster Management GUI, where the active Controller will be highlighted.



Alternatively, you can run the following CLI command:

```
fazbdctl show members
```

The controller appears in the *Role* column

Blade

This refers to the physical blade server enclosed within the FortiAnalyzer-BigData chassis.

The Chassis Management Module

The Chassis Management Module (CMM) is used to remotely manage and monitor server hosts, power supplies, cooling fans, and networking switches for the FortiAnalyzer-BigData unit. The CMM comes with a web management utility that consolidates and simplifies system management for the FortiAnalyzer-BigData chassis.

The web management utility aggregates and displays data from the CMM and provides the following key management features:

- Enables administrators to view in-depth hardware-level status information using a single interface.
- Provides an OS-independent, remote graphical console.
- Allows remote users to power control all or each of the blades.

Columnar Data Store

Unlike the traditional FortiAnalyzer data storage, FortiAnalyzer-BigData relies on the Kudu storage engine, which allows to store data in a columnar fashion.

Tables are split into contiguous segments called tablets, which represent a generic logical unit ready for further replication and parallelization. The replication factor is "3", which means three copies are stored in the system: one original copy and two replicated ones. The replicas are guaranteed to spread across different nodes for fault tolerance.

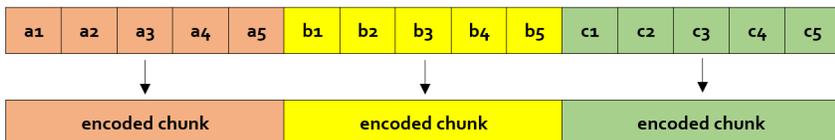
A tablet with N replicas (usually three or five) can continue to accept writes with up to $(N - 1) / 2$.

Kudu uses the Raft consensus algorithm for the election of masters and leader tablets, as well as determining the success or failure of a given write operation, which enforces the data integrity across replicas.

Row layout:



Columnar layout:



This allows aggressive compression and possibility of querying only necessary columns.

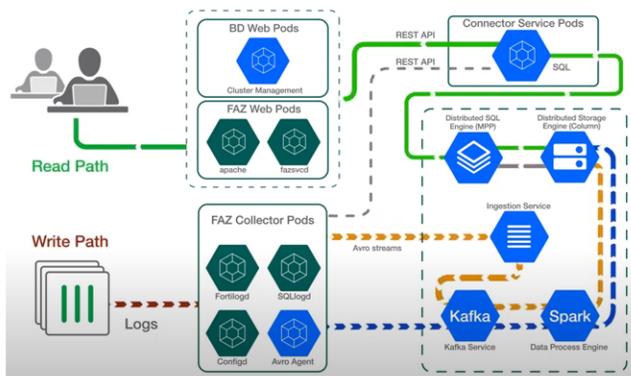
Kudu data store also makes stored log data mutable, which means that stored log events can be changed later.

Controller

This refers to the [Security Event Manager Controller](#).

Data Flow

The following diagram shows the logging write and read path inside the platform:



Write Path:

The write path consists of the following steps:

1. Logs generated from logging devices arrive at the Main Host where the fortilogd process stores them temporarily on a local storage that serves as a buffer before distributing them across all hosts.
2. Logs are packed into a memory-efferent binary format and then streamed by the SQLlogd daemon to the "Ingestion Services", which are Kubernetes Pods processes running on each of Security Manager hosts and

acting as an interface accepting the log data and forwarding it to the Distributed Stateful Workload engines.

3. First Distributed Stateful engine receiving the log data from Ingestion Services is Kafka processes, acting as BD buffering platform for the logs.
4. Spark distributed engine then pulls the logs from Kafka buffers in parallel streams and processes them in fault-tolerant micro-batches. These micro-batches are streamed to Kudu acting as a distribute storage engine. Kudu processes store the logs in a columnar data store, where they can be easily retrieved.

Read Path:

The read path consist of the following steps:

1. An admin tries to access the Logs via *FortiView* or *Log View*.
2. The logs are queried via REST API and Connector Services Pods that consist of Kubernetes Pods processes providing the interface between Main Host and the Security Event Manager hosts.
3. The REST API calls are translated to SQL queries and forwarded to Impala acting as a Distributed SQL Engine.
4. Impala coordinates and distributes the queries across Kudu processes, allowing so called Massively Parallel Processing (MPP).
5. The logs pulled from Kudu are then forwarded to FortiAnalyzer web services and displayed in GUI.

Data Management

The concept of "Archive logs" and "Analytics logs" is not valid for FortiAnalyzer-BigData. All logs are load-balanced across all hosts, where data is compressed, replicated, and available for immediate analytics.

Logs are stored in CFile format with a size of approximately 300 bytes post replication (x3) and compression.

Host

This refers to one of the server hosts in the FortiAnalyzer-BigData system.

Instances

Also known as Service instances. This refers to the instance serving the service. There are usually multiple instances running behind a service load balance.

Main host

The FortiAnalyzer-BigData main host is responsible for collecting logs and providing the services for FortiView, Log View, Reports, and more.

For FortiAnalyzer-BigData units, the main host runs on Blade A1.

Roles

The FortiAnalyzer-BigData hosts are categorized into different roles according to the kind of stateful services running on them. The roles are assigned automatically during the cluster initialization. The placement of those stateful services on each role is designed to achieve optimized performance, high data and service availability and scalability, and is immutable after the cluster is initialized. In a scaling-out scenario (see [Scaling FortiAnalyzer-BigData on page 103](#)), additional hosts can be added as Data nodes to the existing cluster. For

FortiAnalyzer-BigData units, the additional hosts can be added on the extender chassis to the existing cluster in the main chassis.

FortiAnalyzer-BigData has the following roles and services:

- Main Node (this role exists for FortiAnalyzer-BigData-4500G only)
 - Log Collector
 - Main Services
- Master Node
 - Consul
 - Controller Service
 - HDFS Datanode
 - HDFS Journalnode
 - Impala
 - Kafka Broker
 - Kudu Master
 - Kudu Tablet Server
 - Zookeeper
- MetaStore Node
 - HDFS Datanode
 - HDFS Namenode
 - Hive Metastore
 - Impala
 - Impala Catalog
 - Impala Statestore
 - Kafka Broker
 - Kudu Tablet Server
- Data Node
 - HDFS Datanode
 - Impala
 - Kafka Broker
 - Kudu Tablet Server

Services

This refers to the Security Event Manager services that are responsible for security data management, security data processing, storage, cluster management, and more.

Storage Pool

A Storage Pool is a set of one or more ADOMs. Storage pools provide fine-grained control over the data retention policy and improves the query and ingestion performance. Each storage pool can have its own data retention policy that controls the maximum age (in days) and disk utilization of the data. ADOMs within the same storage pool share the storage pool resource.

We recommend grouping ADOMs with similar log rates and data retention requirements into a storage pool. For example, group small ADOMs (in terms of log rate and data volume) into one storage pool and larger ADOMs in

another. If different sized ADOMs are grouped into one storage pool, the query performance on the smaller ADOMs will be affected by the larger ADOMS.

FortiAnalyzer-BigData hardware environment

FortiAnalyzer-BigData 4500F and FortiAnalyzer-BigData 4500G units are available. Each unit is a 4U chassis with two network switch modules and 14 blades in the enclosure.

In the FortiAnalyzer-BigData-4500F unit:

- The first blade is responsible for log collection and services for FortiView, Log View, Reports, and more.
- The remaining 13 blades, also known as hosts, are responsible for the web GUI, log storage, data processing, and analytics.
- The switch modules are 10Gbps.

In the FortiAnalyzer-BigData-4500G unit:

- All blades are used for all functions (log collection, log storage, data processing, analytics, and so on).
- The switch modules are 25Gbps.

Unit	CPU	RAM	Storage
4500F	Blade 1: two 2.1GHz Intel Xeon 12 Core and 24 Thread (12C24T) CPUs Remaining 13 blades: two 2.1GHz Intel Xeon 8 Core 16 Thread (8C16T) CPUs per blade	128GB per blade	Blade 1: two 750GB NVMe SSD Remaining 13 blades: two 7.68TB NVMe SSD per blade
4500G	two 2.3GHz Intel Xeon 16 Core 32 Thread (16C32T) CPUs per blade	196GB per blade	two 7.68TB NVMe SSD per blade

The two network switch modules on the back of the chassis have different functions.

- Switch Module #1 connects to the FortiAnalyzer-BigData cluster's internal network. Use this switch only when you need to scale the existing Security Event Manager by adding new appliances.
- Switch Module #2 is the External Switch Module used to expose the FortiAnalyzer-BigData to external networks.

The Chassis Management Module (CMM) sits between the two switch modules in the middle of the back panel.

For more information about the CMM, see the appropriate FortiAnalyzer-BigData Getting Started Guide for your appliance on the [Fortinet Document Library](#).

Connect to the FortiAnalyzer-BigData CLI

After configuring the FortiAnalyzer-BigData network, you can use the IP addresses to access the FortiAnalyzer-BigData Main CLI or the Security Event Manager Controller and manage the system.

To connect to the FortiAnalyzer-BigData Main CLI:

1. Establish an SSH connection to the *Main Host* IP you configured in the set up process.
For the set up process, see the FortiAnalyzer-BigData Getting Started Guide for your appliance on the [Fortinet Document Library](#).
2. Log in using the administrator credentials you created when setting up the administrator account.
If you did not create a new administrator credential, use the default credentials of username `admin` with no password.

To connect to the Security Event Manager Controller:

1. Establish an SSH connection to the Cluster Management IP you configured in the set up process.
For the set up process, see the FortiAnalyzer-BigData Getting Started Guide for your appliance on the [Fortinet Document Library](#).



If the Cluster Management IP is not reachable, you can SSH to the Main CLI first (see [To connect to the FortiAnalyzer-BigData Main CLI.](#)) and then SSH to the Controller host or any of the cluster hosts using its internal IP. (For example, to SSH to the Controller host, use `exec ssh root@198.18.1.2`).

The IP it is in can be determined by this format: `198.18.{chassis_id}.{blade_id}` where `198.18*` is the default internal subnet.

-
2. Log in using the default username `root` and password `fortinet@123`.
 3. After establishing a connection, you can use the `fazbdct1` CLI commands to manage the cluster. For more information, see the FortiAnalyzer-BigData CLI Reference on the [Fortinet Doc Library](#).



Fortinet strongly recommends that you update the password with the `fazbdct1 set password` command.

GUI overview

FortiAnalyzer-BigData retains the same general GUI as the base FortiAnalyzer. Modules can be accessed from the tree menu on the left-side pane. In addition to the base FortiAnalyzer, there is also a *Cluster Manager* and *Global Search* in FortiAnalyzer-BigData.

Global Search	The Global Search module enables you to perform advanced log search across multiple FortiAnalyzer-BigData clusters. See Global search on page 51 .
Cluster Manager	The Cluster Manager module enables you to manage hosts, services, logs, queries, jobs, and data resources in the Security Event Manager. See Cluster Manager on page 13 .
System Settings	Configure system settings such as network interfaces, administrators, system time, server settings, and others. You can also perform maintenance and firmware operations. See the FortiAnalyzer Administration Guide .

Cluster Manager

The Cluster Manager module enables you to manage hosts, services, logs, queries, jobs, and data resources in the Security Event Manager.

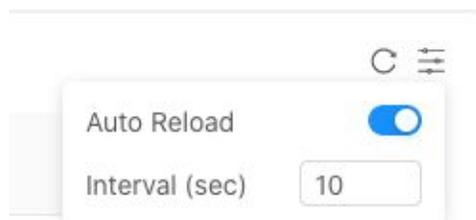
Use the navigation bar to access all the pages within the module.

Section name	Description
Hosts	The Host page enables you to centralize Security Event Manager. It also shows the service assignments as well as resource usage of each host within the Security Event Manager. For more information, see Host management on page 17 .
Services	The Services page enables you to manage the configurations and life cycle of the Security Event Manager. For more information, see Service management on page 21 .
Monitor	The Monitor section contains three pages: <ul style="list-style-type: none">• Dashboard: Provides a customizable visualization for system metrics.• Log and Metrics: Contains an Explorer tool that enables you to search the logs and metrics that FortiAnalyzer-BigData produces.• Health: Provides push notifications for system health checks and other events. For more information, see Monitor on page 26 .
Jobs	The Jobs page manages system jobs and custom jobs.

Section name	Description
	<ul style="list-style-type: none"> System jobs include data retention jobs which removes data outside of the retention period. From this page, you can run jobs, and see the status and history of all your jobs. Custom jobs can be set up with built-in templates or customizable playbooks. For more information, see Job management and automation on page 67.
Data	<p>The Data page enables you to manages the data life cycle of your storage pools as well as data backups and restores.</p> <p>For more information, see Data management on page 75.</p>
Settings	<p>The Settings page includes the following:</p> <ul style="list-style-type: none"> Network: View and edit the blade network. Upgrade: Upgrade the FortiAnalyzer-BigData. For more information, see Upgrade FortiAnalyzer-BigData on page 99. System: View the system version and collect diagnostic logs from the Security Event Manager services and hosts.

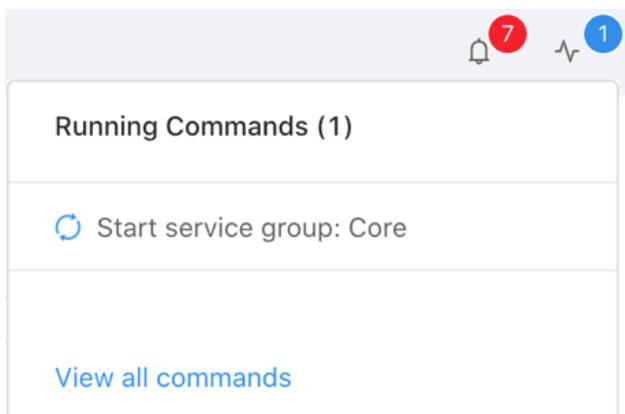
Custom refresh settings

When viewing tables in the Cluster Manager, you can manually refresh the data in a table by clicking *Refresh* , or you can set up an automatic reload timer. Click *Custom Settings*  at the top-right corner of a table to configure the refresh setting.



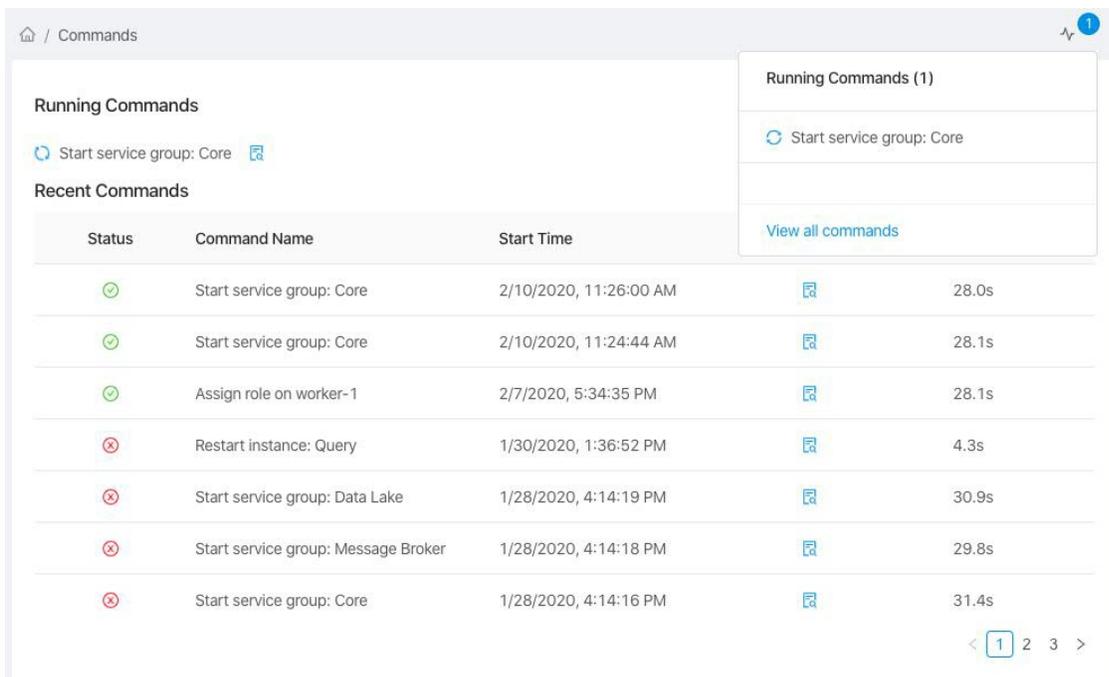
Commands management

There is a Commands icon  in the top-right corner of each page that notifies you whenever a command is running in the background. You can click the icon to expand the Commands snapshot view and see all currently running commands.



To access the Commands Manager page

1. Click *Commands*  in the top-right of each page.
The Commands snapshot view loads, showing all the currently running commands.
2. Click *View all commands* at the bottom of the snapshot to view the full list of commands.

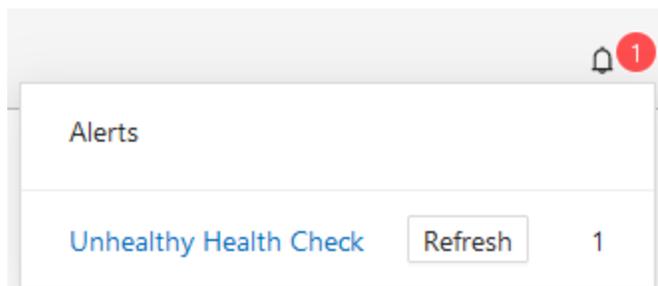


3. The icon by each command indicates if the command was executed successfully.

-  The command was successfully executed.
-  The command failed.

Notifications management

There is a *Notifications* icon  in the top-right corner of each page that notifies you each time there is a notification. You can click the icon to expand the alerts snapshot view and see more details. Clicking an item directs you to the page related to the alert.



For the specific alerts such as the "Unhealthy Health Check" alert, you can click the *Refresh* button to refresh all information related to that check.

Host management

The Host page has a table that provides an overview of all the hosts in the Security Event Manager. You can use the *Actions* column to manage hosts.

Status	Host	Role Type	Instances	CPU Usage	Memory Usage	Disk Usage	Actions
<input type="checkbox"/>	blade-198-18-1-2	Master Node	▶ 7 Instance(s)	36.6%	82.8 GB / 125.6 GB	12.1 TB / 14.1 TB	Status Details
<input type="checkbox"/>	blade-198-18-1-3	Master Node	▶ 7 Instance(s)	31.2%	77.2 GB / 125.6 GB	9.2 TB / 14.1 TB	Status Details
<input type="checkbox"/>	blade-198-18-1-4	Master Node	▶ 7 Instance(s)	32.2%	82.3 GB / 125.6 GB	11.1 TB / 14.1 TB	Status Details
<input type="checkbox"/>	blade-198-18-1-5	Data Node	▶ 3 Instance(s)	36.5%	80.1 GB / 125.6 GB	10.3 TB / 14.1 TB	Status Details
<input type="checkbox"/>	blade-198-18-1-6	Data Node	▶ 3 Instance(s)	37.9%	84.6 GB / 125.6 GB	10.1 TB / 14.1 TB	Status Details
<input type="checkbox"/>	blade-198-18-1-7	Data Node	▶ 3 Instance(s)	33.7%	87 GB / 125.6 GB	12.3 TB / 14.1 TB	Status Details
<input type="checkbox"/>	blade-198-18-1-8	Data Node	▶ 3 Instance(s)	34.4%	79.3 GB / 125.6 GB	10.3 TB / 14.1 TB	Status Details
<input type="checkbox"/>	blade-198-18-1-9	Data Node	▶ 3 Instance(s)	38.5%	80.1 GB / 125.6 GB	11 TB / 14.1 TB	Status Details
<input type="checkbox"/>	blade-198-18-1-10	MetaStore Node	▶ 7 Instance(s)	29.5%	87.9 GB / 125.6 GB	10.7 TB / 14.1 TB	Status Details
<input type="checkbox"/>	blade-198-18-1-11	MetaStore Node	▶ 7 Instance(s)	28.0%	84.1 GB / 125.6 GB	10.1 TB / 14.1 TB	Status Details
<input type="checkbox"/>	blade-198-18-1-12	Data Node	▶ 3 Instance(s)	26.4%	80.9 GB / 125.6 GB	9.3 TB / 14.1 TB	Status Details
<input type="checkbox"/>	blade-198-18-1-13	Data Node	▶ 3 Instance(s)	34.3%	89.9 GB / 125.6 GB	10.8 TB / 14.1 TB	Status Details

The Host table contains the following columns:

Column header	Description
Status	<p>There are three icons that represent the status of the host:</p> <ul style="list-style-type: none"> ● The host is healthy. ● The host is in poor health. ● A command is currently running on the host.
Host	The name of the host.
Role Type	<p>Each host can have one role. For more information about each role, see Roles on page 9.</p> <ul style="list-style-type: none"> ● Main Node (for FortiAnalyzer-BigData-4500G only) ● Master Node ● MetaStore Node

Column header	Description
	<ul style="list-style-type: none"> • Data Node • Unassigned: The host is new and does not have an assigned role. Click <i>new</i> to assign a role to that host (see Role assignment on page 19).
Instances	The number of Service instances on each host. You can expand the row to see which instances are on each host and their current status.
CPU Usage	The percentage of the CPU being used.
Memory Usage	How much memory is being used.
Disk Usage	How much space is being used on a disk.
Actions	<p>You can perform the following actions on each host:</p> <ul style="list-style-type: none"> • Show Details: Show the status details. • Restart: Restart the host. • CLI Console: Open the CLI console for the host.

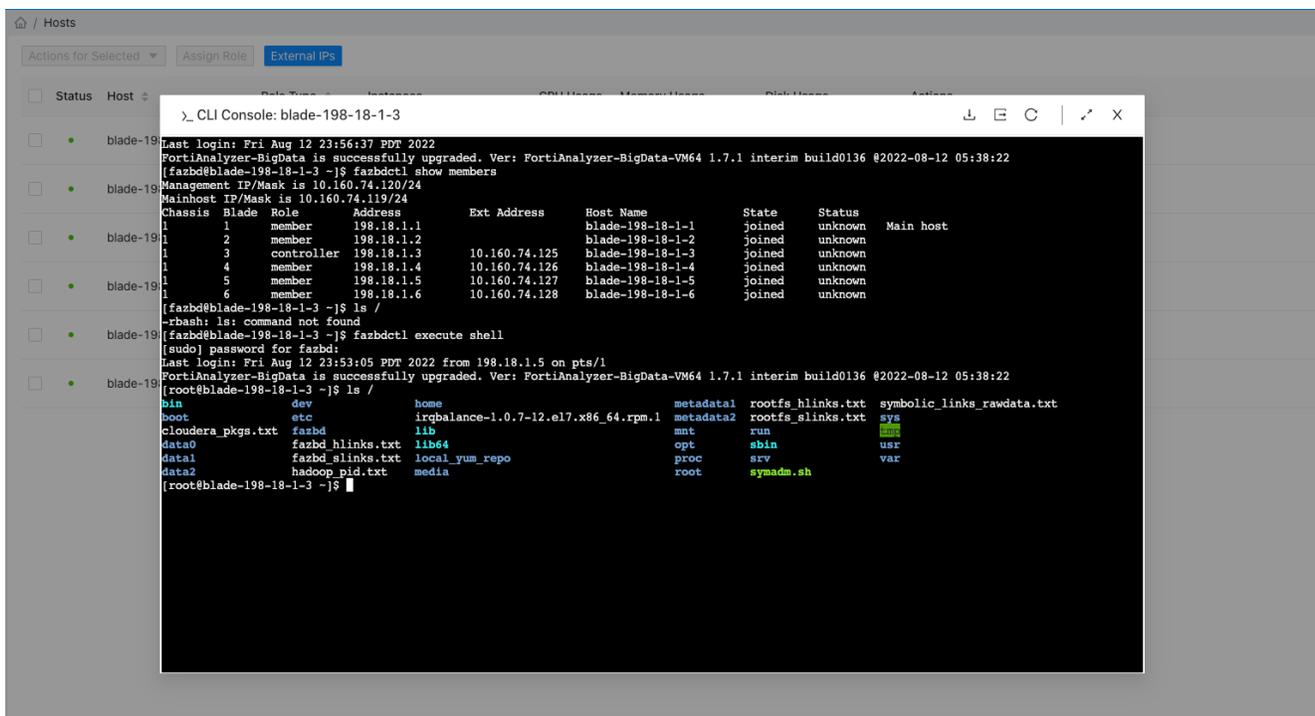
CLI Remote Console

You can remotely enter the CLI console of a FortiAnalyzer-BigData cluster host on *Cluster Manager* web GUI. To enter the CLI console of a host, click the *CLI Console* in the menu under the *Actions* column. Each host has its own CLI console and multiple consoles can be opened at the same time.

The screenshot shows the 'Hosts' page in the Cluster Manager web GUI. At the top, there are navigation elements: a home icon, the text '/ Hosts', and buttons for 'Actions for Selected', 'Assign Role', and 'External IPs'. Below this is a table with columns: Status, Host, Role Type, Instances, CPU Usage, Memory Usage, Disk Usage, and Actions. Three hosts are listed:

Status	Host	Role Type	Instances	CPU Usage	Memory Usage	Disk Usage	Actions
●	blade-198-18-1-3	Master Node	▶ 9 Instance(s)	10.2%	7.3 GB / 15.6 GB	50 GB / 1 TB	Status Details ⋮
●	blade-198-18-1-1	Main Node	▶ 2 Instance(s)	7.8%	7.3 GB / 19.6 GB	18.9 GB / 1 TB	Status Details ⋮
●	blade-198-18-1-2	Master Node	▶ 9 Instance(s)	13.0%	8.3 GB / 15.6 GB	48.4 GB / 1 TB	Status Details ⋮

A context menu is open over the 'Actions' column of the second host (blade-198-18-1-1), showing two options: 'Restart' and 'CLI Console'.



The CLI Console view supports the following actions:

Button	Action	Description
	Download History	Download the current console output.
	Open as new tab	The CLI console will open in a new browser tab.
	Reconnect	Reconnect to the host's CLI console.
	Maximize	Click to expand the console to fit the window. Click again to restore the window size.

By default, the CLI console has limited access to the shell that allows only some of the `fazbdctl` and `fazbdadm` commands. You can enter the secure shell with the following command:

```
fazbdctl execute shell
```

FortiAnalyzer-BigData cluster host's OS password will be prompted. For more information, see *Set Cluster Hosts OS password* and *Enable Secure Shell* in the [FortiAnalyzer-BigData CLI Reference Guide](#).

Role assignment

Hosts that have an Unassigned role type are flagged with a *new* notification.

<input type="checkbox"/>	Status ▾	Host Name ▾		Role Type ▾	Address ▾	Instances ▾	Actions
<input type="checkbox"/>	●	blade-10-0-2-39	new	Unassigned	10.0.2.39	0 Instance(s)	Assign Role

You can assign a role to a host by clicking *Assign Role* in the Actions column.

To assign a role to a host

1. In the Actions column, click *Assign Role*.
The *Assign Role dialog* loads.
2. Select the role you want to assign to the host.



At this time, you can only assign the Data Node role.

3. Click *Assign* to confirm your selection.
FortiAnalyzer-BigData begins the role assignment process.

Service management

The Services page has a table with information about all the services running on your system. This table provides an overview of all your services and enables you to monitor and manage all the services running in the system.

Status	Service Group	Services	Actions
●	Core	10	Start Stop Restart
●	Message Broker	2	Start Stop Restart
●	Data Lake	3	Start Stop Restart
●	Metastore	4	Start Stop Restart
●	Monitor	9	Start Stop Restart

The Services table contains the following columns:

Column header	Description
Status	There are five icons that represent the status of the host: <ul style="list-style-type: none">● The services are healthy.● A command is currently running.● There is a problem with the service.● Services within the service group are experiencing issues.● The service has stopped.
Service Group	Service Groups are a way to group and categorize individual services. Click on the Service Group to access the Service Configuration page and manage the services contained inside. By default, FortiAnalyzer-BigData has four pre-defined Service Groups (see Service groups on page 22).
Services	The number of services running in each group.
Actions	There are three actions you can perform on each Service Group or service. <ul style="list-style-type: none">• Start: Start the service group or a specific service.• Stop: Stop the service group or a specific service.• Restart: Restart the service group or a specific service.

Service groups

Services are organized into Service groups, which can contain several services. Each service can further contain multiple instances running on a host. By default, FortiAnalyzer-BigData has four pre-defined Service groups that contain the following services:

Service Group	Services within the Service group
Core	<ul style="list-style-type: none">• Catalog• Query• Ingestion• Hyperscale IPFIX Ingestion• Pipeline• Controller Service• Controller Failover• Management Portal• Management Server• Management Task
Message Broker	<ul style="list-style-type: none">• Kafka• Rabbitmq
Data Lake	<ul style="list-style-type: none">• Impala• Kudu• HDFS
Metastore	<ul style="list-style-type: none">• Zookeeper• Consul• Redis• Stolon
Monitor	<ul style="list-style-type: none">• Monitor Portal• Metrics Server• Metrics Exporters• Log Server

Service details

To access the Service Details page, click the name of the Service group.

The screenshot shows the 'Data Lake' service group page. It has two tabs: 'Instances' (selected) and 'Configuration'. Below the tabs is a dropdown menu for 'Actions for Selected' and a search icon. The main content is a table with the following data:

Status	Service Name	Instances	Pending Configs	Actions
<input type="checkbox"/> ●	Impala	19	0	Start Stop Restart
<input type="checkbox"/> ●	Kudu	16	0	Start Stop Restart
<input type="checkbox"/> ●	HDFS	18	0	Start Stop Restart

The Service Details page contains all the services grouped under the Service group. You can expand each service to see the instances it contains, and manually start, stop, or restart those services.

The screenshot shows the expanded details for the 'Kafka' service. It features a table with the following data:

Status	Instance Name	Instance State	Host Name	Address	Actions
<input type="checkbox"/> ●	Kafka Broker	Started	● blade-10-0-2-2	10.0.2.2	Start Stop Restart
<input type="checkbox"/> ●	Kafka Broker	Started	● blade-10-0-2-32	10.0.2.32	Start Stop Restart
<input type="checkbox"/> ●	Kafka Broker	Started	● blade-10-0-2-33	10.0.2.33	Start Stop Restart

Some services may contain configurations that you can modify via the Configuration tab.

Home / Services / Message Broker

Instances **Configuration**

Kafka Broker

Configurations have been saved, and 1 configurations are pending. Go to [Instances](#) tab to apply. ×

Kafka Broker

* Number of Threads for Disk I/O:
num.io.threads

Data Directories:
log.dirs

* Data Retention Time: Hours
log.retention.hours

* Data Retention Size: Bytes ⓘ
log.retention.bytes

* Enable Auto Creation of Topic:
auto.create.topics.enable

* Number of Partitions:
num.partitions

* Default Replication Factors:
default.replication.factor

Kafka Heap Options:

[Reset to Default](#) [Reset to Last Applied](#) [Save](#)

To modify service configurations



The FortiAnalyzer-BigData default configurations are optimized for performance, availability, and scalability. Configure these settings with caution as improper configurations can have a negative impact on the entire system, and even lead to system failure or data loss. Approach these options with great care and when in doubt, err on the side of caution.

1. From the Service page, click the Service group name to access the Service Configuration page.
2. Click *Configuration* to switch to the Configuration tab.
3. Modify the fields as needed.
4. Once you are finished, click *Save*.
Once you save the changes, you must apply the changes.



You can click *Reset to Default* to reset the changes to the default configurations, or click *Reset to Last Applied* to reset the configurations to the last changes you applied.

5. To apply the configurations, click *Instances* to return to the Instance tab. The number in the Pending Configs column changes to reflect the number of configurations that are pending.
6. In the Actions column, click *Apply Config* to apply the changes.

Monitor

From the Navigation bar, you can expand the Monitor section to access three pages:

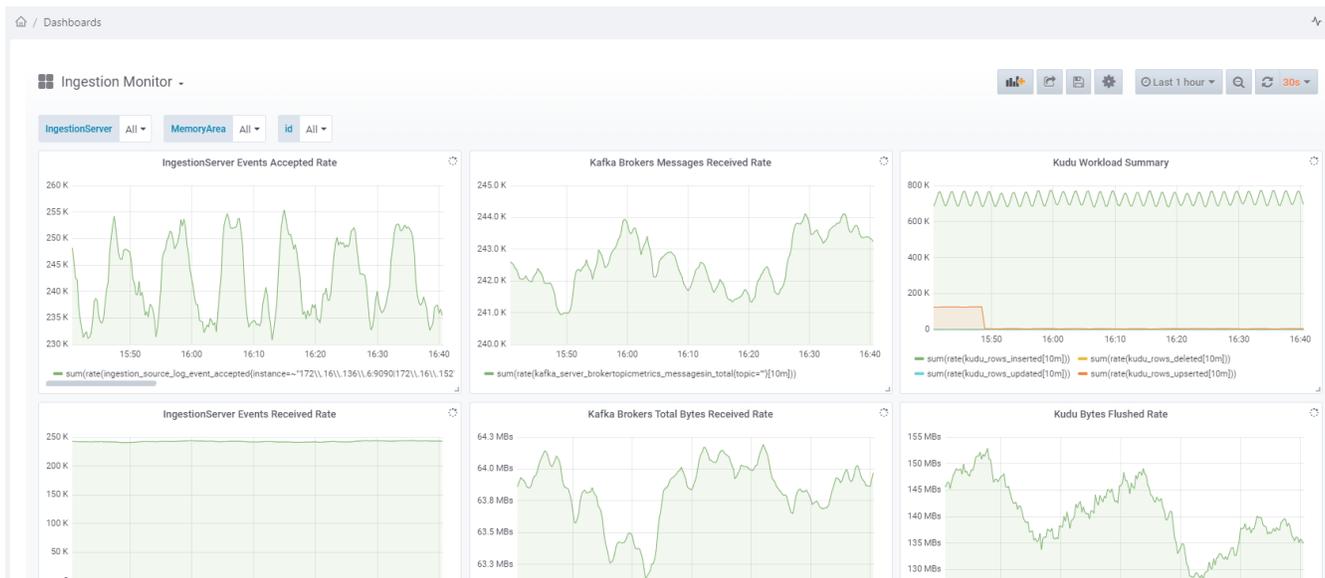
- [Dashboards](#)
- [Logs and metrics](#)
- [Health](#)

Dashboards

The Dashboards page displays both real-time monitoring and historical trends of your system metrics.

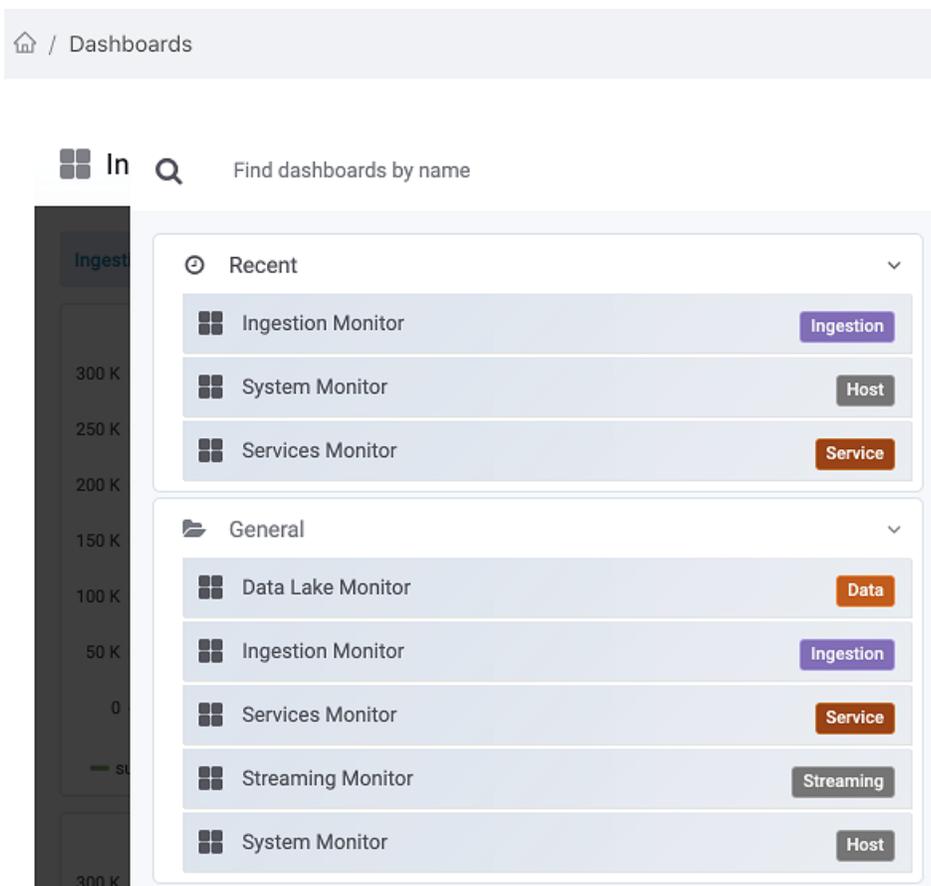
From the Dashboard, you can:

- Select specific data to focus on and filter your results to narrow down your view
- Customize the panels on your dashboard
- Use panels to see more information and set alerts



Filtering the Dashboard

You can filter the dashboard to focus on different areas of focus. By default, the Dashboard shows statistics relating to data ingestion. You view built-in dashboards by clicking the title of each page (for example, Ingestion Monitor) and selecting a topic from the drop-down list.



You can view the following built-in dashboards:

- Ingestion
- Data Lake
- Services
- Streaming
- System

In some views, you can filter your results to show information from a specific server, node, memory area, ID, and more. You can also narrow down results to a specific time period and set the refresh rate.

Customizing the Dashboard

You can customize the FortiAnalyzer-BigData dashboard by adding new panels, creating custom settings, and saving those settings. Once you've customized the dashboard, you can share the dashboard.

Dashboard actions	Description
Add panel	Add a panel to your dashboard. Once a blank panel appears on the dashboard, you can select the following actions: <ul style="list-style-type: none"> • Add Query: Choose what metrics to track, • Choose Visualization: Choose how you want to visualize the data. • Convert to row: Convert a group of panels into a collapsible row.
Share Dashboard	Share the dashboard with a link or by exporting a JSON file.
Save Dashboard	Save all the changes you've made to the dashboard.
Dashboard settings	
General	Configure general settings for the current dashboard. The FortiAnalyzer-BigData dashboard is built on Grafana. For more information about using dashboard features, refer to the official Grafana documentation .
Annotations	Add annotations to mark points on a graph.
Variables	Add variables to change the data being displayed in the dashboard.
Links	Add a link to your dashboard so you can go to other dashboards and websites directly.
Versions	See the revision version history for the dashboard.
JSON Model	See the JSON model that defines the dashboard.

Using panels

The Dashboard contains panels that display specific metrics. You can drag and drop each panel to rearrange your Dashboard view, or stretch the panel to see more details. Click the drop-down menu on each panel to get a list of available actions.

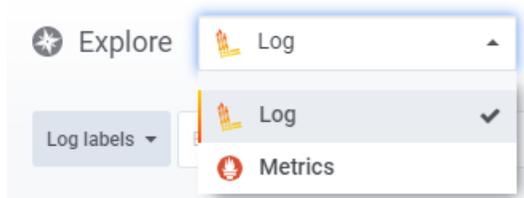


Panel menu actions	Description
View	Enlarge the panel to see a more detailed view of the graph.
Edit	You can customize the panel to show specific queries, change the way you visualize data, and set alerts rules to inform you when certain conditions are met.
Share	There are two ways to share your panels: <ul style="list-style-type: none"> • Create a direct link to the particular panel. • Create a snapshot of the panel with sensitive data stripped out.
Explore	View the historical logs and metrics for the panel.
More	
Duplicate	Add a duplicate of the panel to your dashboard.
Copy	Create a copy of the pane. You can paste the panel to the Dashboard from <i>Add panel</i> .
Panel JSON	See the JSON model that defines the panel.
Export CSV	Export a CSV file with panel data.
Toggle Legend	Click to display or conceal the panel legend.
Remove	Remove the panel from the Dashboard.

Logs and metrics

The Logs & Metrics page contains all the logs and metrics that FortiAnalyzer-BigData produces.

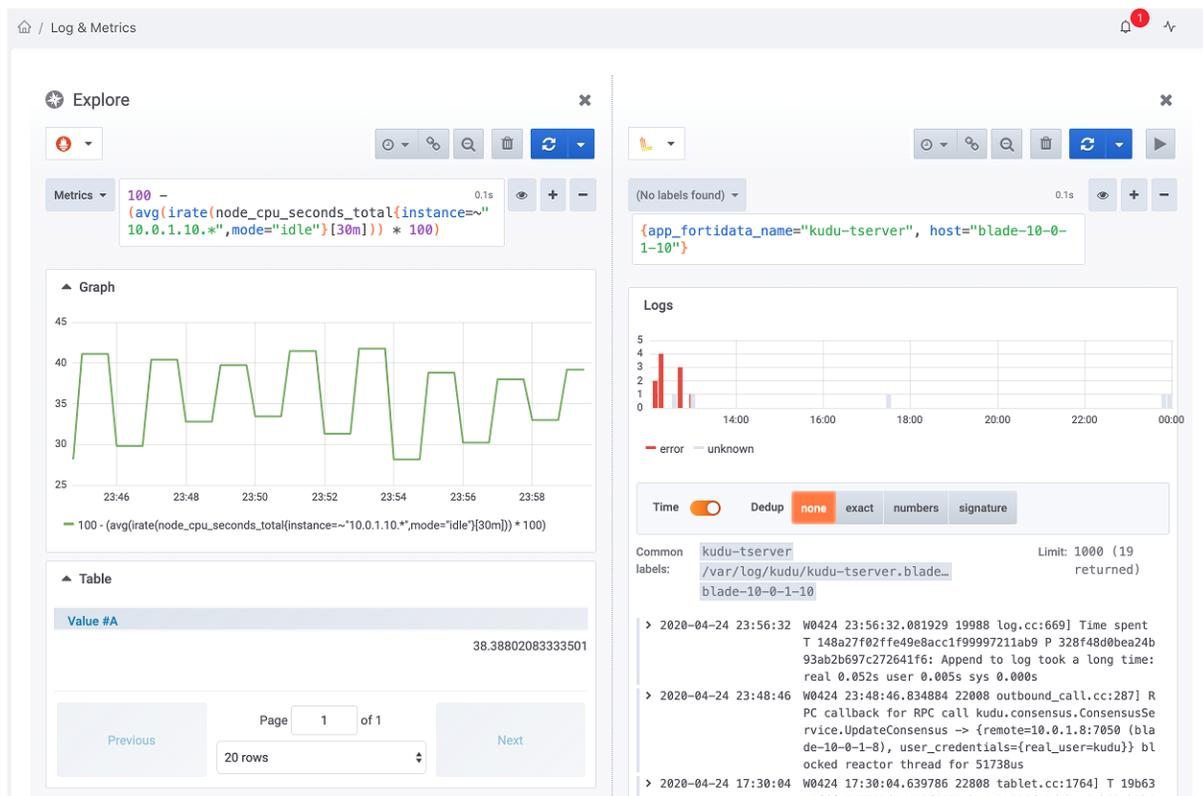
You can use the Explore search tool to switch between the Logs or Metrics view. The default selection is Logs.



- Logs are immutable records of discrete events that happened over time in the system.
- Metrics are a set of numbers that give information about a particular process or activity.

After you select a view, you can search for the particular log or metric that you want to see. You can add filters to show results from a certain time range.

The Logs and Metrics page has a Split screen feature which enables you to compared two different Logs or Metrics at the same time. Click *Split* to create a side-by-side comparison view.



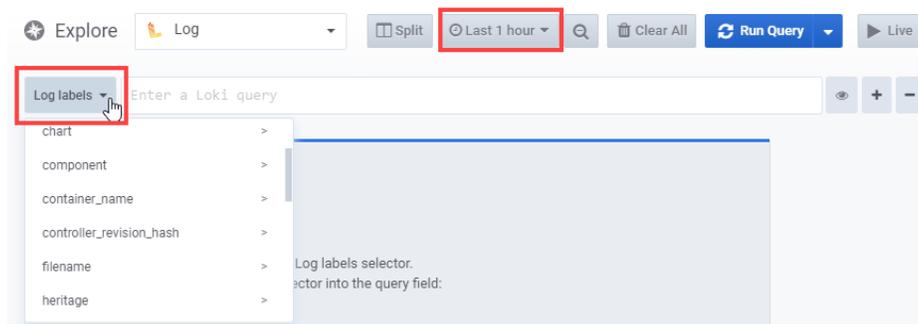
Explore logs

A log query has two main components:

- a log stream selector; and
- a search expression.

Choosing a log stream

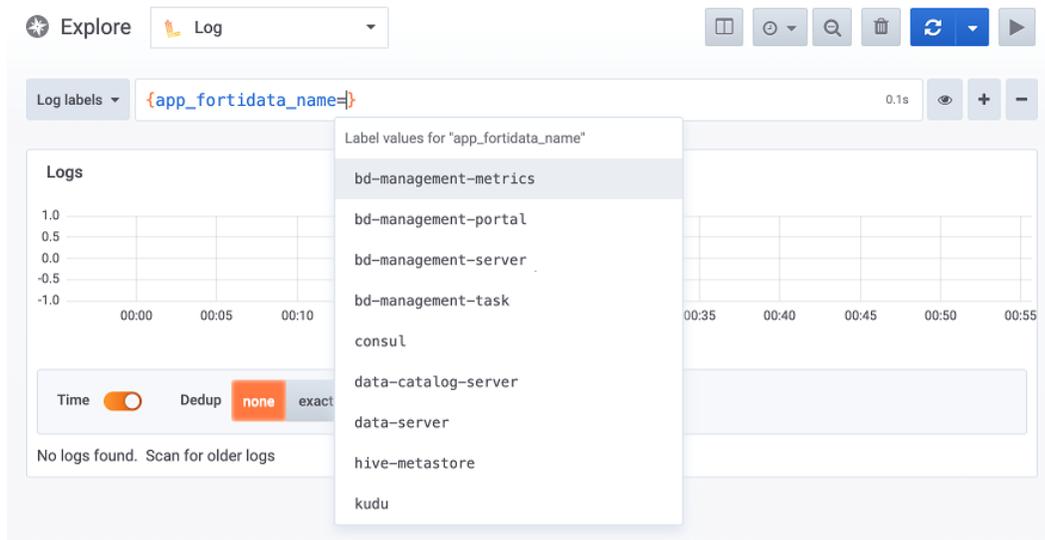
Choose a log stream by clicking the *Log labels* button next to the search bar, and select from the available log streams in your time range (the default time range is Last 1 hour).



If there are no logs in the selected time range, the log label of the log will not show up in the label list.

Entering a search expression

You can start a search query by using the search field's autocomplete feature. Enter a curly bracket { in the search field to see a suggested list of labels. You can browse through the suggested labels with your cursor or arrow keys and press the Tab key to select a label. Press the Enter key to execute the query.



The log stream selector is wrapped inside curly braces {} with the key and value of selecting labels. You can select multiple labels by using commas, for example:

```
{app_fortidata_name="ingestion-server", host="blade-10-0-1-10"}
```

This example selects the ingestion-server log on host blade-10-0-1-10.

After you choose a selector, you can follow up by entering a search expression to filter the results further. Search expressions can be in a text or regex expression, for example:

```
{app_fortidata_name="data-server"} |= "ERROR"
{app_fortidata_name="ingestion-server"} |~ "Starting.*engine"
{host="blade-10-0-1-10"} != "INFO"
```

You can chain the operators in order to search the log lines and satisfy all filters. For example:

```
{app_fortidata_name="ingestion-server"} |= "ERROR" != "timeout"
```

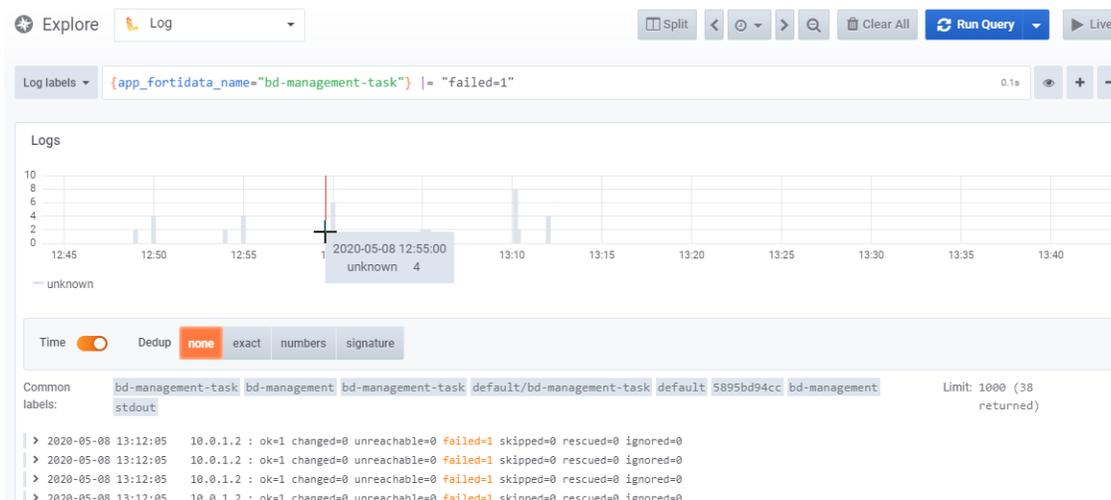
Supported operators:

- |= line contains a string.
- != line does not contain a string.
- |~ line matches regular expression.
- !~ line does not match regular expression.

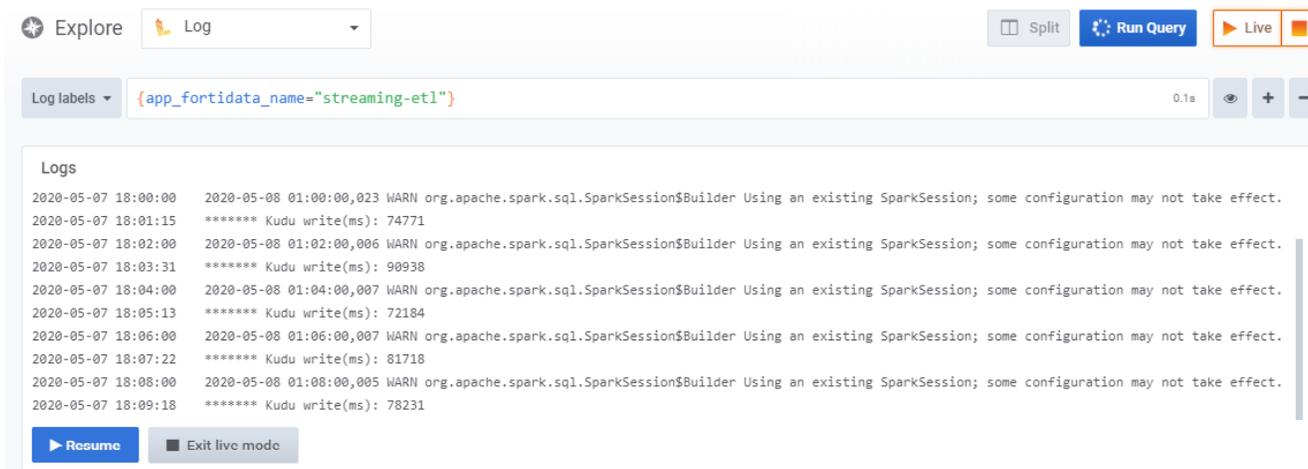
For more details, refer to the Loki query language (LogQL) documentation.

Log query results

After you run a query, search results are presented as either a list of log rows and/or a bar graph. For results with a bar graph, the time is placed on the x-axis while log count is on the y-axis. You can click and drag on the bar chart to narrow down the time range.



You can also click the *Live* button to enter Live Tailing mode and see logs changes in real-time.



If you use a search expression, you can see the context for each filtered result by hovering your mouse over a result and clicking the *Show Context* link by each result.

Time Dedup **none** exact numbers signature

Common labels: bd-management-task bd-management bd-management-task default/bd-management-task default 5895bd94cc bd-management stdout Limit: 1000 (38 returned)

```

> 2020-05-08 13:12:05 10.0.1.2 : ok=1 changed=0 unreachable=0 failed=1 skipped=0 rescued=0 ignored=0 Show context
> 2020-05-08 13:12:05 10.0.1.2 : ok=1 changed=0 unreachable=0 failed=1 skipped=0 rescued=0 ignored=0
> 2020-05-08 13:12:05 10.0.1.2 : ok=1 changed=0 unreachable=0 failed=1 skipped=0 rescued=0 ignored=0
> 2020-05-08 13:12:05 10.0.1.2 : ok=1 changed=0 unreachable=0 failed=1 skipped=0 rescued=0 ignored=0
> 2020-05-08 13:10:27 10.0.1.2 : ok=6 changed=2 unreachable=0 failed=1 skipped=0 rescued=0 ignored=0
> 2020-05-08 13:10:27 10.0.1.2 : ok=6 changed=2 unreachable=0 failed=1 skipped=0 rescued=0 ignored=0
> 2020-05-08 13:10:15 localhost : ok=0 changed=0 unreachable=0 failed=1 skipped=0 rescued=0 ignored=0
> 2020-05-08 13:10:15 localhost : ok=0 changed=0 unreachable=0 failed=1 skipped=0 rescued=0 ignored=0
> 2020-05-08 13:10:14 localhost : ok=0 changed=0 unreachable=0 failed=1 skipped=0 rescued=0 ignored=0
> 2020-05-08 13:10:14 localhost : ok=0 changed=0 unreachable=0 failed=1 skipped=0 rescued=0 ignored=0
> 2020-05-08 13:10:14 localhost : ok=0 changed=0 unreachable=0 failed=1 skipped=0 rescued=0 ignored=0
> 2020-05-08 13:10:14 localhost : ok=0 changed=0 unreachable=0 failed=1 skipped=0 rescued=0 ignored=0
> 2020-05-08 13:10:14 localhost : ok=0 changed=0 unreachable=0 failed=1 skipped=0 rescued=0 ignored=0
> 2020-05-08 13:10:13 localhost : ok=0 changed=0 unreachable=0 failed=1 skipped=0 rescued=0 ignored=0
> 2020-05-08 13:10:13 localhost : ok=0 changed=0 unreachable=0 failed=1 skipped=0 rescued=0 ignored=0
> 2020-05-08 13:05:23 10.0.1.2 : ok=6 changed=3 unreachable=0 failed=1 skipped=0 rescued=0 ignored=0
    
```

When you click *Show Context*, a new window loads enabling you to see the context of that particular result.

Found 10 rows. Load 10 more

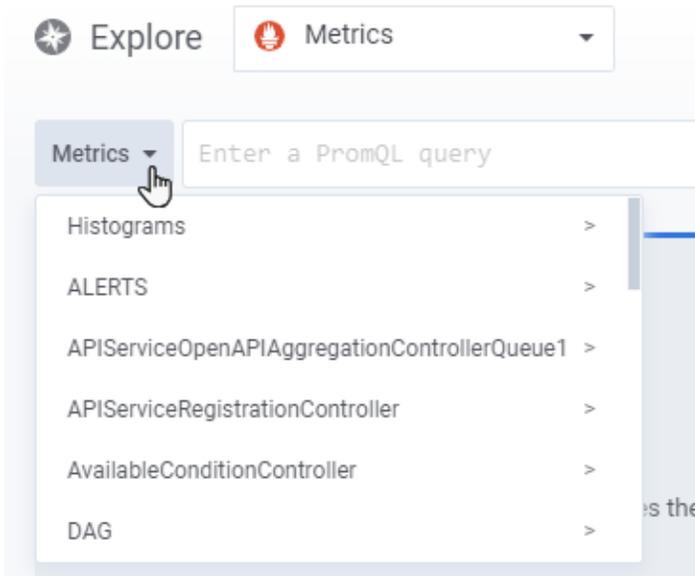
```

get process information of datanode ----- 0.63s
get available space for each data dir ----- 0.73s
get data dir content ----- 1.24s
run command to get datanode usage info ----- 2.56s
Gathering Facts ----- 3.48s
=====
Saturday 25 April 2020 07:49:36 +0000 (0:00:00.030) 0:00:11.021 *****
blade-10-0-1-10 : ok=23 changed=7 unreachable=0 failed=1 skipped=2 rescued=0 ignored=0
Hide context
PLAY RECAP *****
fatal: [blade-10-0-1-10]: FAILED! => {"changed": false, "msg": "This datanode free space is below thresholds"}
Saturday 25 April 2020 07:49:36 +0000 (0:00:00.059) 0:00:10.991 *****
TASK [check if node's free space is below thresholds] *****
ok: [blade-10-0-1-10]
Saturday 25 April 2020 07:49:36 +0000 (0:00:00.070) 0:00:10.932 *****
Found 10 rows. Load 10 more
    
```

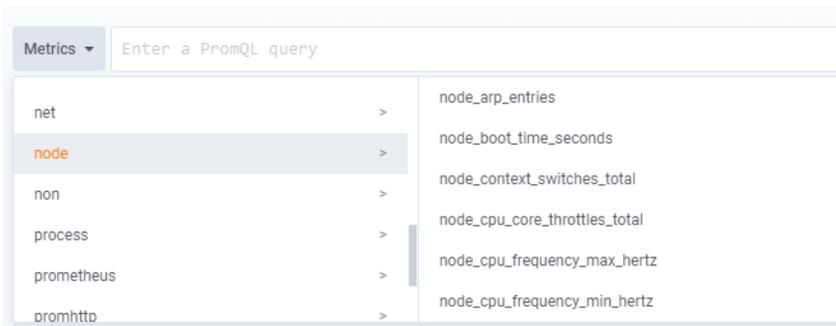
Explore metrics

Access the Explore Metrics view by changing the Explore field selection to *Metrics*.

To search for a metrics, click the *Metrics* dropdown to open a hierarchical menu with available metrics.

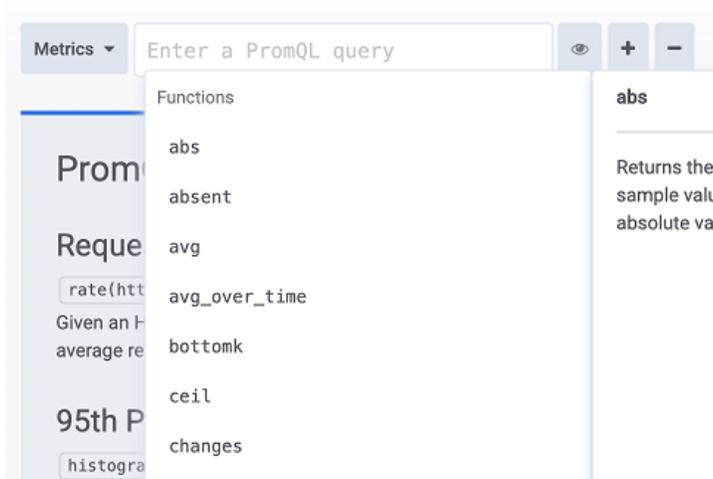


Metrics are grouped by prefixes, for example, all Node metrics are grouped under the "node" prefix.



After you select a metrics key, the data is represented with a graph and table. The raw data is listed in the table with label keys as columns and the label values and metric values as rows.

You can also start a query by pressing the **Ctrl** key in search box to display suggestions for metric names and functions. Press the **Enter** key to execute.

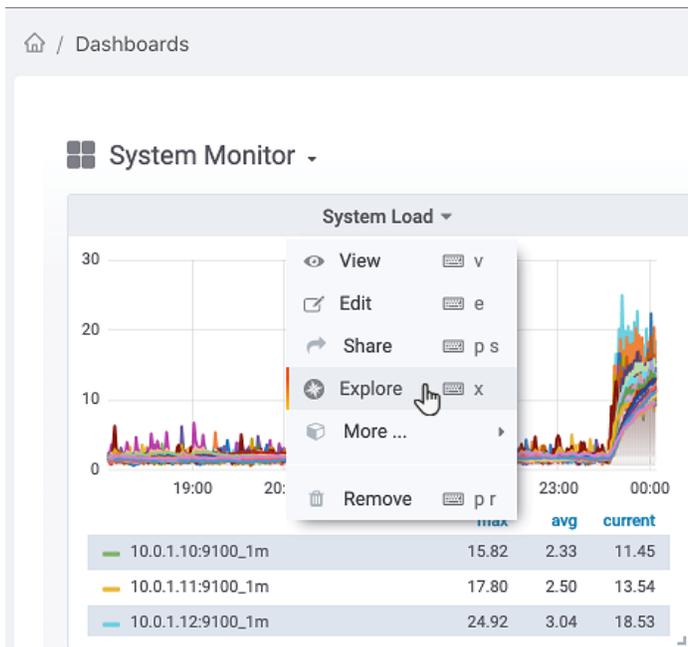


For more details, refer to the Prometheus Query Language documentation.

Accessing a specific metrics from the Dashboard

You can also access a specific metric by drilling down from a dashboard panel.

Find the specific panel you want to see metrics data for, click the panel title and select *Explore*.



Health

In the Health page, you can set alerts for system health checks, and configure how you want to receive your alerts.

Health Check

The Health Check tab displays a table containing all predefined health checks in the system.

Status	Health Check	Schedule	Last Test Result	Last Updated	Actions
Failed	HDFS DataNode Health Check	At 22 minutes past the hour	Test Result	4/25/2020, 12:49:36 AM	Run Test Configure
Success	HDFS Health Check	At 36 minutes past the hour	Test Result	4/25/2020, 12:36:18 AM	Run Test Configure
Success	HDFS JournalNode Health Check	At 51 minutes past the hour	Test Result	4/25/2020, 12:51:11 AM	Run Test Configure
Success	HDFS NameNode Health Check	At 37 minutes past the hour	Test Result	4/25/2020, 12:37:27 AM	Run Test Configure

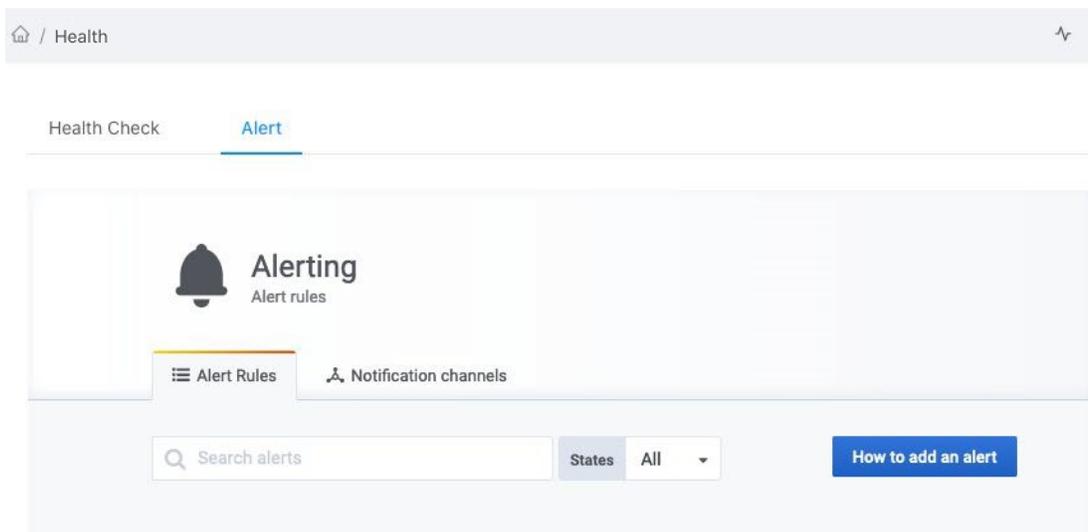
The Health Check table contains the following columns:

Column header	Description
Status	Indicates if the health check was a success or failure. If a health check fails, you can click <i>Expand</i> in the item row to see the error message.
Health Test	Shows what health check was run. You can click the name to see the history for that health check. FortiAnalyzer-BigData only saves the last 300 records for each health check.
Schedule	Shows how often the health check is run.
Last Test Result	View the full health test result by clicking <i>Test Result</i> .
Last Updated	The last time the health test was run.
Actions	You can perform two actions on the health test:

Column header	Description
	<ul style="list-style-type: none"> • <i>Run Test</i>: Manually start the health test. • <i>Configure</i>: Change how often the test is run by configuring the scheduling settings.

Alert

The Alert tab enables you to search through your existing alerts and set rules on how you receive alerts. You can also configure how you want to receive push notifications through various notification channels such as email, Slack, PagerDuty, WebHook, and more.



Notification channel alerts

You can add new ways of receiving alerts by adding a channel and specifying the channel type.

To create a notification channel with email

The following example shows how to set up the SMTP server and create an email notification channel.

1. Go to *Services > Monitor > Configuration*.
2. Enable *SMTP*.
3. In the *SMTP Host* field, enter the SMTP server address and SMTP port. The format is <SMTP server address>:<SMTP port number>. For example, `smtp.gmail.com:587`.
4. In the *SMTP TLS Policy* field, select *TLS policy*.
5. (Optional) Enable *SMTP authentication* if authentication is required.
6. In the *SMTP Auth User* and *SMTP Auth Password* fields, enter the username and password for authentication.
7. Click *Save*.

8. Go to the *Instances* tab and click *Apply Config*. The configuration changes take effect and triggers the `Enable Smtpt` command.
Wait for the command to finish running.
9. Go to *Monitor > Health > Alert > Notification channels* and click *New channel*.
10. In the *Name* field, enter a name for the channel.
11. In the *Type* field, select *Email*.
12. In the *Addresses* field, enter the destination email addresses for notifications. Separate multiple email addresses with a semi-colon (;).
13. Click *Test* and check if you can receive the test alert email.
14. Click *Save* once you have verified the email channel alert works.

To create a notification channel with Slack Incoming Webhook

The following example shows how to create a notification channel with Slack Incoming Webhook and set up an alert.

1. Go to *Monitor > Health > Alert > Notification channels* and click *Add channel*.
2. In the *Name* field, enter a name for the channel.
3. In the *Type* field, select *Slack*.
4. You can choose how you want to configure your alert.
In this example, enable the *Include image* toggle so a snapshot of your Slack chart can be sent with the alert.
5. In the *URL* field, enter your Slack Incoming Webhook URL.
For instructions on how to create a Slack Incoming Hook, refer to the Slack documentation.
6. In the *Token* field, enter the in the Slack "Bot User OAuth Access Token" in order to allow the generated image to be uploaded via Slack's file.upload API method.
7. In Slack, invite the bot to the channel you want to send notifications to and add the Slack channel name to the *Recipient* field.
8. Click *Send Test* and check if you can see the test message in your Slack channel with the Webhook hooked.

9. Once you have verified that the channel alert works, click Save.

The screenshot shows the 'Edit Notification Channel' interface. At the top, there's a breadcrumb 'Health / Health' and tabs for 'Health Check' and 'Alert'. The main title is 'Edit Notification Channel'. Below it, there are several configuration rows:

- Name:** Slack
- Type:** Slack
- Default (send on all alerts):**
- Include image:**
- Disable Resolve Message:**
- Send reminders:**

Below these are 'Slack settings':

- Url:** https://hooks.slack.com/services/xxxxxxxxxxxxx...
- Recipient:** #alerts
- Username:**
- Icon emoji:**
- Icon URL:**
- Mention:**
- Token:** xoxb-0000000xxxxxxxxxxxxxxxx000000000000xx

At the bottom, there are four buttons: Save (green), Send Test (blue), Delete (red), and Back (grey).

Custom alert rules

You can create custom alert rules from Dashboard panels and have it sent to a specified notification channel.

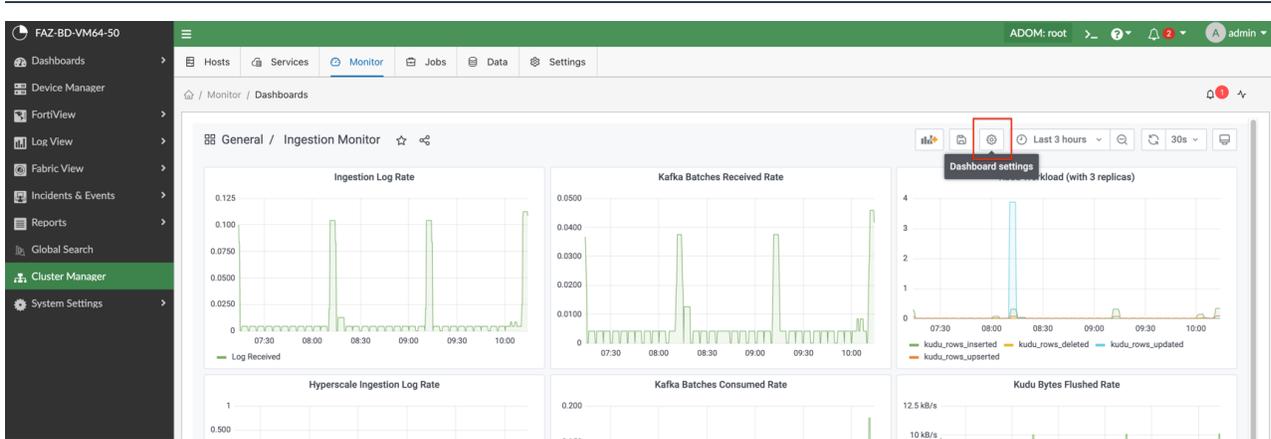
To create a custom alert for a notification channel:

The following example shows how to create custom alert rule that can be sent directly to the example Slack notification channel.

1. Go to *Monitor > Dashboards* and select a panel for which you want to create an alert.
2. Click *Dashboard Settings*.



Fortinet recommends creating alerts on custom dashboards. If alerts are set up on the default dashboard, they may be lost when the default dashboard is updated during FortiAnalyzer-BigData firmware upgrades.



3. Click Save As.

The screenshot shows the FortiAnalyzer interface with the following elements:

- Header:** FAZ-BD-VM64-50, FortiView, and navigation tabs: Hosts, Services, Monitor (selected), Jobs, Data, Settings.
- Left Sidebar:** Dashboards, Device Manager, FortiView, Log View, Fabric View, Incidents & Events, Reports, Global Search, Cluster Manager (highlighted), System Settings.
- Breadcrumbs:** / Monitor / Dashboards
- Page Title:** Ingestion Monitor / Settings
- General Tab:**
 - Name:** Ingestion Monitor
 - Description:** (empty field)
 - Tags:** Ingestion (tag), New tag (enter key to add) Add
 - Folder:** General
 - Editable:** Editable (selected), Read-only
 - Time options:**
 - Timezone:** Default
 - Auto refresh:** 5s,10s,30s,1m,5m,15m,30m,1h,2h,1d
 - Now delay now:** 0m
 - Hide time picker:** (disabled)
 - Panel options:**
 - Graph tooltip:** Default
- Buttons:** Save dashboard, Save As... (highlighted), Delete Dashboard

4. Enter a *Dashboard name* and click *Save*.

Save dashboard as...
×

Dashboard name

Folder

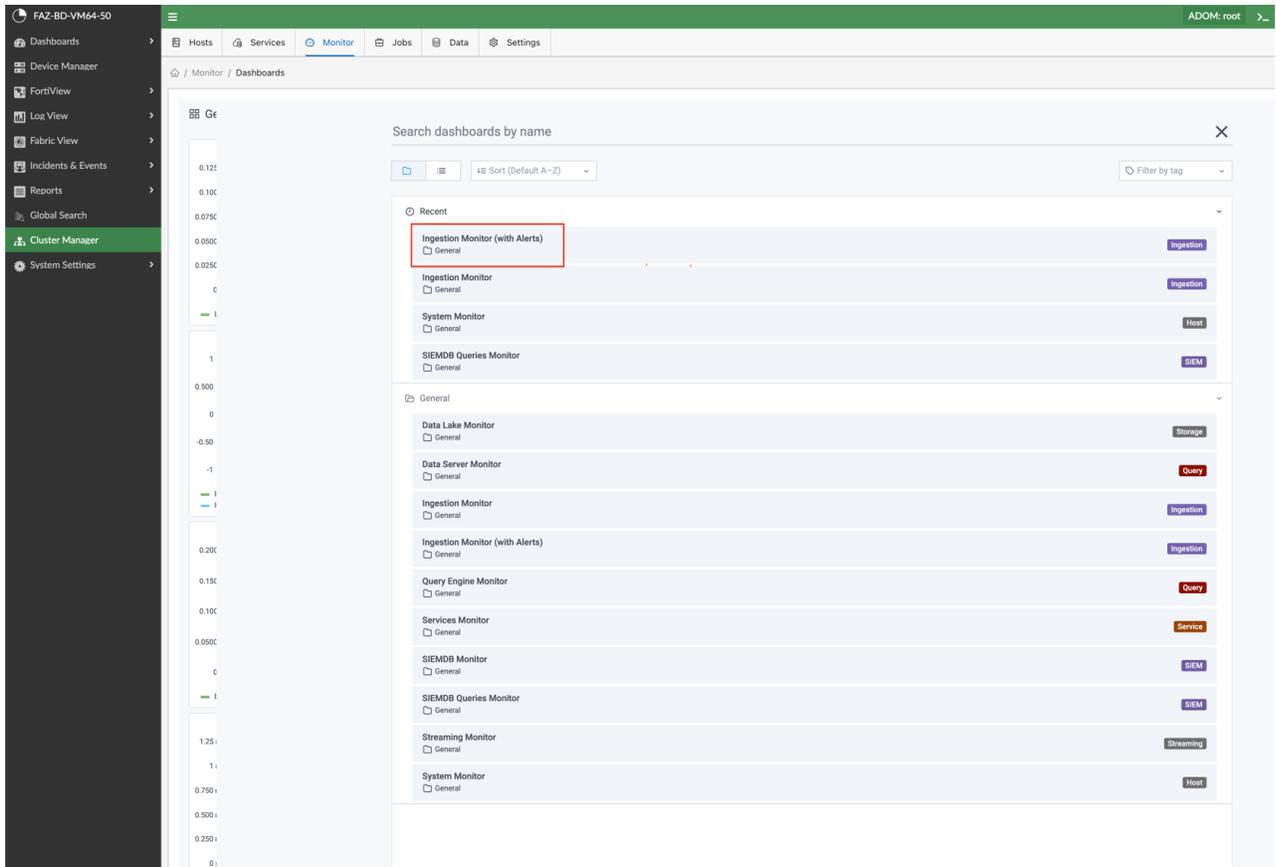
General
▾

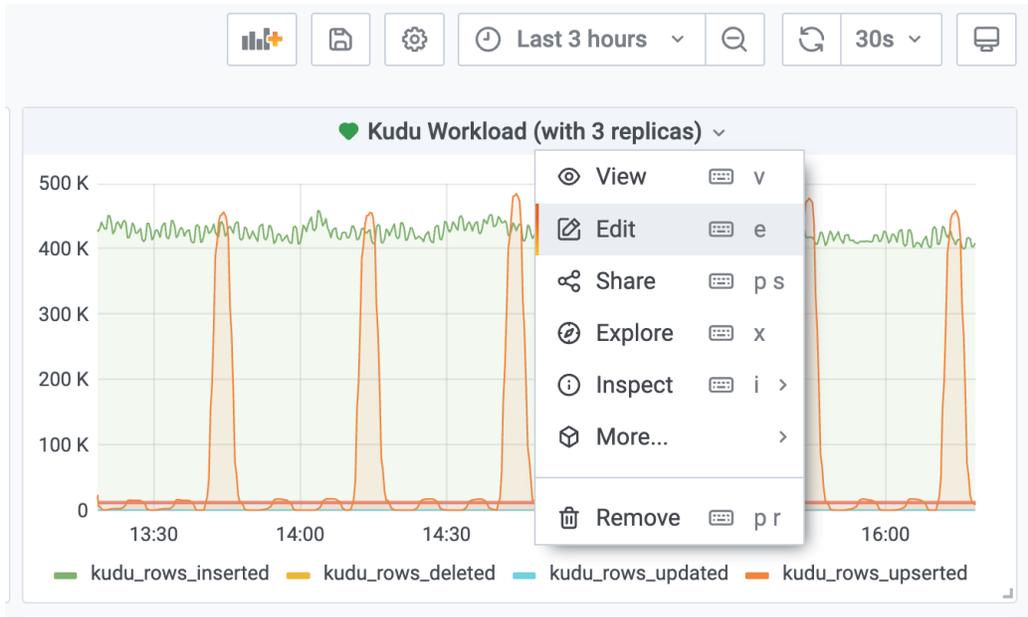
Copy tags

Save

Cancel

5. Navigate to the created Dashboard, click a panel title, and click *Edit*.

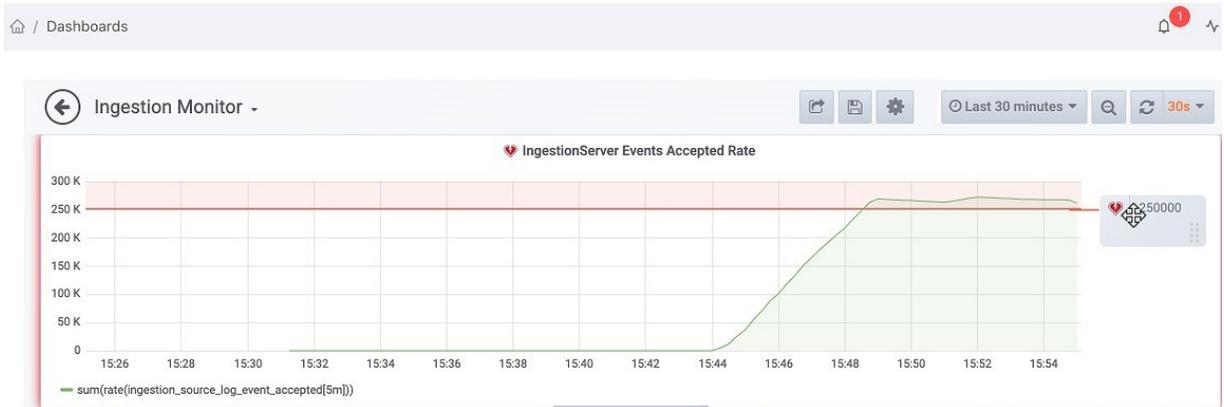




The panel's detailed view displays.

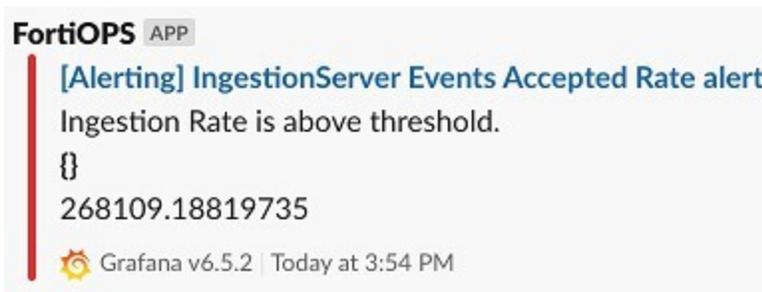
6. Click Alert (bell icon) to access the *Alert* view.
7. Click *Create Alert* to specify conditions that trigger the alert.
8. You can create conditions through two different methods:
 - By making queries in the *Conditions* section.

- By dragging the threshold bar in the graph to indicate an allowable threshold level.



9. After you have defined your condition, select the *Notification Channel* and click *Test Rule* to test the alert rule.
10. Click *Save* to save your settings.

If your conditions are configured correctly, you should receive an alert with snapshot resembling the following:



Hyperscale firewall logging support

You can configure FortiAnalyzer-BigData as a log server for a Fortinet Hyperscale firewall that supports hardware logging with Hyperscale SPU log offload feature. FortiAnalyzer-BigData can collect, store, and query log messages sent as NetFlow v10, which is compatible with IP Flow Information Export (IPFIX) format, or Syslog over UDP from a Hyperscale firewall. For more information, see [Hyperscale Firewall Hardware logging](#) in the Fortinet Doc Library.

Set up Security Manager hosts external IP addresses

When receiving log messages from Hyperscale FortiGate, each Security Manager host of FortiAnalyzer-BigData can be exposed as a distributed log collector to distribute the log traffic load. A set of external IP addresses for the hosts that your FortiGate can reach are required to receive the log message traffic.

To set up external IP addresses for Security Manager hosts:

1. Go to *Cluster Manager > Settings > Network > Blade Network* and click *Edit*.

Blade Network

[Edit](#)

Name	Internal IP	Internal Interface	External IP	External Interface(s)
blade-198-18-1-2	198.18.1.2	Network adaptor 1	10.106.2.52	Network adaptor 2
blade-198-18-1-3	198.18.1.3	Network adaptor 1	10.106.2.53	Network adaptor 2
blade-198-18-1-4	198.18.1.4	Network adaptor 1	10.106.2.54	Network adaptor 2
blade-198-18-1-5	198.18.1.5	Network adaptor 1	10.106.2.55	Network adaptor 2
blade-198-18-1-6	198.18.1.6	Network adaptor 1	10.106.2.56	Network adaptor 2
blade-198-18-1-7	198.18.1.7	Network adaptor 1	10.106.2.57	Network adaptor 2

2. Set the *Default Gateway* and *Netmask* to your internal network.

3. Enter an IP address for each host.

The screenshot shows a dialog box titled "Set Blade External IPs" with a close button (X) in the top right corner. At the top, there are two tabs: "Set by host" (which is selected and highlighted in blue) and "Set with range". Below the tabs, there are several input fields:

- Gateway: 10.106.2.254
- Netmask: 255.255.255.0
- Interface: Network adaptor 2 (with a dropdown arrow)
- blade-198-18-1-2: 10.106.2.52
- blade-198-18-1-3: 10.106.2.53
- blade-198-18-1-4: 10.106.2.54
- blade-198-18-1-5: 10.106.2.55
- blade-198-18-1-6: 10.106.2.56
- blade-198-18-1-7: 10.106.2.57

 At the bottom right, there are "Cancel" and "Apply" buttons.

4. If you have a range of continuous IP addresses, you can click *Set with range* and specify a *Start External IP* to automatically increment and set the IP addresses to all hosts.

The screenshot shows the same "Set Blade External IPs" dialog box, but now the "Set with range" tab is selected and highlighted in blue. The "Set by host" tab is now disabled. The input fields are:

- Start External IP: xxx.xxx.xxx.xxx
- Netmask: 255.255.255.0
- Gateway: 10.106.2.254
- Interface: Network adaptor 2 (with a dropdown arrow)

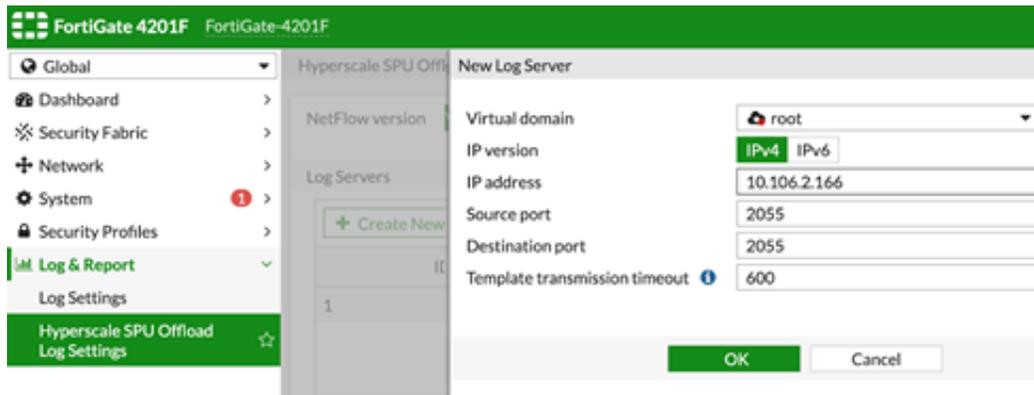
 The "Cancel" and "Apply" buttons are still present at the bottom right.

Configure FortiAnalyzer-BigData as log server on hyperscale FortiGate

After external IP addresses for Security Manager hosts are set, you can configure a FortiGate with Hyperscale firewall features to send NetFlow v10 (IPFIX) or Syslog log messages over UDP to FortiAnalyzer-BigData. For more information, see [Hyperscale Firewall Hardware logging](#) in the [Fortinet Doc Library](#).

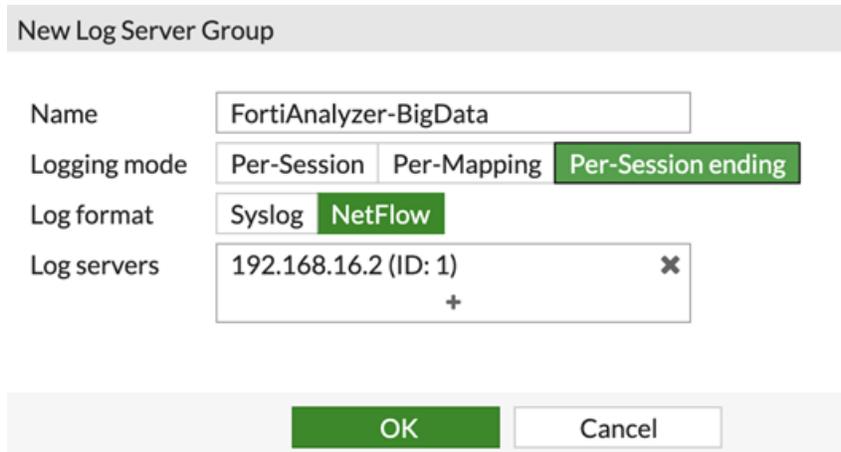
To configure FortiAnalyzer-BigData as NetFlow log server on Hyberscale FortiGate:

1. Go to *Log & Report > Hyperscale SPU Offload Log Settings*.
2. Select *NetFlow version V10*.
3. In *Log Servers*, click *Create New* to add each external IP address of FortiAnalyzer-BigData Security Manager Host.
4. In the *Source port* and *Destination port*, enter 2055.



5. In *Log Servers Groups*, click *Create New* to create a log group.
6. For *Logging mode*, select *Per-Session ending*.
7. For *Log format*, select *NetFlow*.
8. For *Log servers*, add all the log servers created in the previous step.
9. Click *OK*.

The FortiGate is configured to send NetFlow log messages to FortiAnalyzer-BigData.



To configure FortiAnalyzer-BigData as Syslog log server on a hyperscale FortiGate:

1. Go to *Log & Report > Hyperscale SPU Offload Log Settings*.
2. In *Log Servers*, click *Create New* to add each external IP address of FortiAnalyzer-BigData Security Manager Host.
3. In the *Source port* and *Destination port*, enter 514.
4. In *Log Servers Groups*, click *Create New* to create a log group.
5. Set *Logging mode* to *Per-Session ending*.

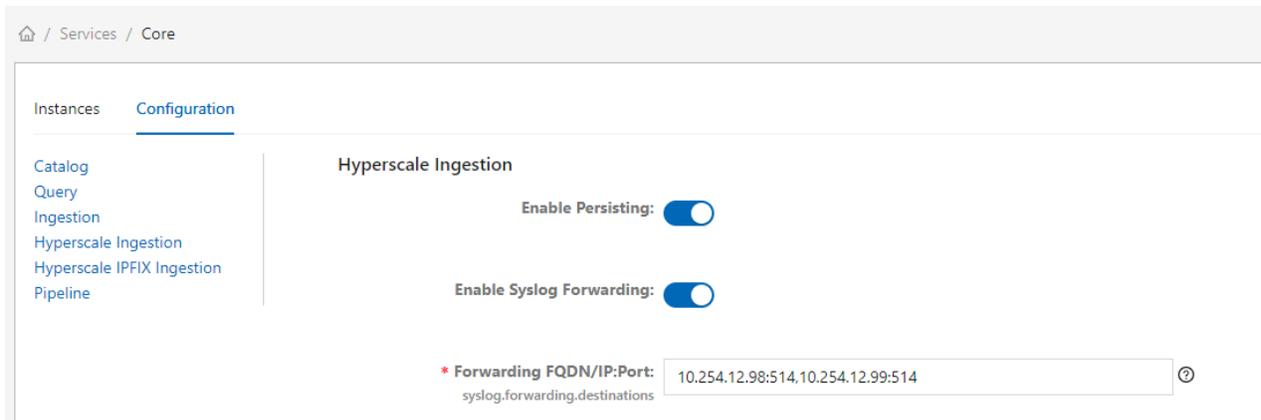
6. Set *Log format* to *Syslog*.
7. For *Log servers*, add all the log servers created in the previous step.
8. Click *OK*. The FortiGate is configured to send Syslog log messages to FortiAnalyzer-BigData.

Configure FortiAnalyzer-BigData Hyperscale Ingestion

By default, the Hyperscale ingestion is disabled in the Security Event Manager. You can also configure FortiAnalyzer-BigData to forward Hyperscale Syslog messages to an external Syslog UDP server:

To configure Hyperscale Ingestion:

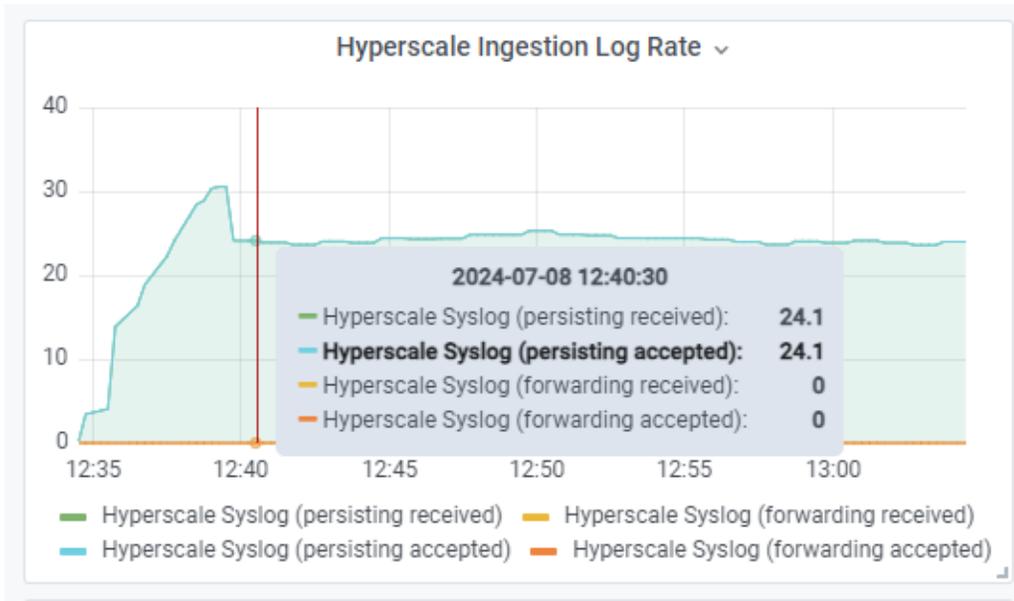
1. Go to *Cluster Manager > Services > Core*.
2. Click the *Configuration* tab and go to the *Hyperscale Ingestion* section.
3. Toggle *Enable Persisting* to enable or disable persisting the logs into the Security Event Manager. This is disabled by default.
4. Toggle *Enable Syslog Forwarding* to enable or disable forwarding to an external Syslog UDP server. This is disabled by default.
5. When *Enable Syslog Forwarding* is enabled, configure the destination Syslog UDP server's IP and port. Multiple servers can be added, separated by a comma.
6. Click *Save* and follow the prompts to apply the configuration.



Device Manager and log rate

When FortiAnalyzer-BigData starts to receive hyperscale log messages, the device appears in Device Manager along with information such as log status, VDOM, log rate, and so on. Click the number to display a graph of historical average log rates of the device.

To view the overall log rate of normal logs and hyperscale logs in different log format, go to *Cluster Manager > Monitor > Dashboards* to view the *Hyperscale Ingestion Log Rate*.



Search Hyperscale log in Log View

If you have Hyperscale logs collected, you can go to *FortiGate > Hyperscale* to view the logs in *Log View*. You can apply filters, click to view log details, and use other features as you search other logs.

The screenshot shows the FortiGate Log View interface. The top navigation bar includes 'Log View', user information 'ADOM: root', and 'admin'. The left sidebar shows a navigation menu with 'Hyperscale' selected. The main content area displays a table of logs for 'All FortiGate' from 'Last 1 Hour' (16:23:25 To 17:23:24). The table has the following columns: #, Date/Time, Device ID, Virtual Domain, NAT Type, Source Address, Post NAT Source Address, and Destination Address. The logs listed are as follows:

#	Date/Time	Device ID	Virtual Domain	NAT Type	Source Address	Post NAT Source Address	Destination Address
1	17:22:00	FGVM0100...	prod-hw1		1.1.1.48	1.2.8.136	
2	17:22:00	FGVM0100...	prod-hw1	4	1.1.3.46	1.1.3.46	2.1.2.98
3	17:22:00	FGVM0100...	prod-hw1	4	1.1.3.139	1.1.3.139	2.1.2.24
4	17:22:00	FGVM0100...	prod-hw1	4	1.1.3.31	1.1.3.31	2.1.2.162
5	17:22:00	FGVM0100...	prod-hw1	4	1.1.2.82	1.1.2.82	2.1.2.245
6	17:22:00	FGVM0100...	prod-hw1	4	1.1.4.98	1.1.4.98	2.1.2.140
7	17:22:00	FGVM0100...	prod-hw1	4	1.1.3.253	1.1.3.253	2.1.1.183
8	17:22:00	FGVM0100...	prod-hw1	4	1.1.4.188	1.1.4.188	2.1.2.19
9	17:22:00	FGVM0100...	prod-hw1	4	1.1.1.54	1.1.1.54	2.1.1.48
10	17:22:00	FGVM0100...	prod-hw1	4	1.1.3.135	1.1.3.135	2.1.1.75
11	17:22:00	FGVM0100...	prod-hw1	4	1.1.1.105	1.1.1.105	2.1.2.199
12	17:22:00	FGVM0100...	prod-hw1	4	1.1.4.10	1.1.4.10	2.1.1.126
13	17:22:00	FGVM0100...	prod-hw1	4	1.1.4.226	1.1.4.226	2.1.2.109
14	17:22:00	FGVM0100...	prod-hw1	4	1.1.1.117	1.1.1.117	2.1.2.12
15	17:22:00	FGVM0100...	prod-hw1	4	1.1.3.200	1.1.3.200	2.1.1.46
16	17:22:00	FGVM0100...	prod-hw1	4	1.1.4.231	1.1.4.231	2.1.1.185
17	17:22:00	FGVM0100...	prod-hw1		1.1.1.48	1.2.8.25	
18	17:22:00	FGVM0100...	prod-hw1		1.1.1.49	1.2.8.172	
19	17:22:00	FGVM0100...	prod-hw1		1.1.1.57	1.2.8.94	
20	17:22:00	FGVM0100...	prod-hw1	4	1.1.1.48	1.1.1.48	2.1.1.48
21	17:22:00	FGVM0100...	prod-hw1	4	1.1.4.225	1.1.4.225	2.1.2.95
22	17:22:00	FGVM0100...	prod-hw1	4	1.1.2.209	1.1.2.209	2.1.1.254
23	17:22:00	FGVM0100...	prod-hw1	4	1.1.3.42	1.1.3.42	2.1.1.220
24	17:22:00	FGVM0100...	prod-hw1	4	1.1.1.160	1.1.1.160	2.1.2.221
25	17:22:00	FGVM0100...	prod-hw1		1.1.1.50	1.2.8.22	
26	17:22:00	FGVM0100...	prod-hw1		1.1.1.55	1.2.8.195	
27	17:22:00	FGVM0100...	prod-hw1	4	1.1.2.33	1.1.2.33	2.1.1.192
28	17:22:00	FGVM0100...	prod-hw1	4	1.1.3.247	1.1.3.247	2.1.2.224
29	17:22:00	FGVM0100...	prod-hw1	4	1.1.3.33	1.1.3.33	2.1.2.47
30	17:22:00	FGVM0100...	prod-hw1	4	1.1.1.95	1.1.1.95	2.1.2.152
31	17:22:00	FGVM0100...	prod-hw1	4	1.1.3.207	1.1.3.207	2.1.2.228
32	17:22:00	FGVM0100...	prod-hw1	4	1.1.4.138	1.1.4.138	2.1.2.13
33	17:22:00	FGVM0100...	prod-hw1	4	1.1.2.212	1.1.2.212	2.1.1.249
34	17:22:00	FGVM0100...	prod-hw1	4	1.1.2.143	1.1.2.143	2.1.2.19
35	17:22:00	FGVM0100...	prod-hw1	4	1.1.2.19	1.1.2.19	2.1.2.59
36	17:22:00	FGVM0100...	prod-hw1		1.1.1.50	1.2.8.111	
37	17:22:00	FGVM0100...	prod-hw1	4	1.1.3.37	1.1.3.37	2.1.2.161
38	17:22:00	FGVM0100...	prod-hw1	4	1.1.4.220	1.1.4.220	2.1.2.161
39	17:22:00	FGVM0100...	prod-hw1	4	1.1.2.230	1.1.2.230	2.1.1.203
40	17:22:00	FGVM0100...	prod-hw1		1.1.1.50	1.2.8.158	
41	17:22:00	FGVM0100...	prod-hw1	4	1.1.2.152	1.1.2.152	2.1.1.183
42	17:22:00	FGVM0100...	prod-hw1		1.1.1.52	1.2.8.147	
43	17:22:00	FGVM0100...	prod-hw1	4	1.1.3.48	1.1.3.48	2.1.1.242
44	17:22:00	FGVM0100...	prod-hw1	4	1.1.2.154	1.1.2.154	2.1.1.136
45	17:22:00	FGVM0100...	prod-hw1	4	1.1.4.94	1.1.4.94	2.1.2.98

At the bottom of the table, there is a pagination control showing '50' items per page, a page number '1' (highlighted), and a refresh button with a timer set to '4.319 Seconds'.

Global search

Global Search lets you explore log messages collected by FortiAnalyzer-BigData across all ADOMs. When searching with a Federation, you can search across multiple clusters.

Use *Global Search* to identify trends in the data with the *Histogram* and detailed log messages at the same time. You can quickly explore log messages by selecting the type and labels and pivoting directly from the fields in the log details with just a few clicks. Perform advanced queries with rich *LogQL* (log query language). Cross-cluster search federation allows you to run searches against one or more remote FortiAnalyzer-BigData clusters and compare the results in a single view.

This section contains the following topics:

- [Starting a global search on page 51](#)
- [Log types \(Global Search\) on page 52](#)
- [Create a new Search Federation \(Example\) on page 58](#)
- [Log Query Language \(LogQL\) on page 64](#)

Starting a global search

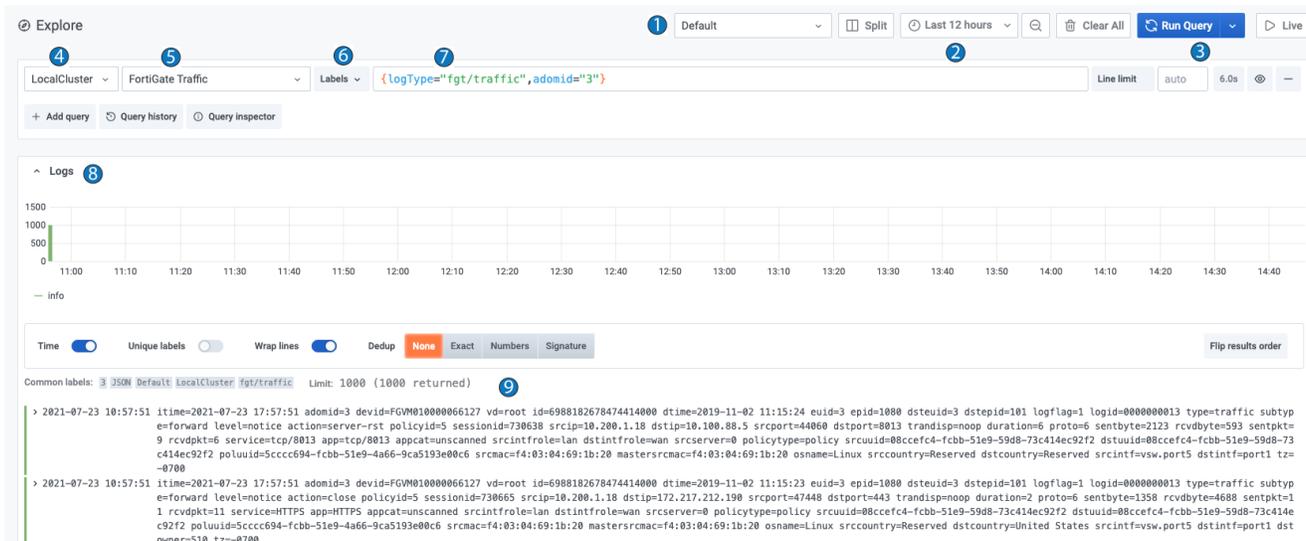
To perform a simple global search, use the default values for the search federation, cluster and log time, then select a log type.

To perform a simple search:

1. Select a search federation. *Default* is the default.
2. Select a cluster. *LocalCluster* is the default.
3. Select the log time. *Last 1 hour* is the default.
4. Select the log type. For example, *FortiGate Traffic*.
5. (Optional) Select the log labels filter.
6. Click *Run Query* to start a search.

Global Search settings

The following image and the corresponding table provide information about each of the Global Search settings.



Search setting	Description
1 Search Federation dropdown	Click the dropdown menu to select a federation. The default is <i>Default</i> .
2 Time Range selection	Click the dropdown menu to select a time range. The default is <i>Last 1 hour</i> .
3 Run Query	Click <i>Run Query</i> to start the search.
4 Server	Click the dropdown to select a server. The default is <i>Local Cluster</i> .
5 Log Type	Click the dropdown to select a log type.
6 Log Label	Click the dropdown to select a log label.
7 Log Query Input	Use this field to enter the log query.
8 Histogram	Displays the log time-range.
9 Log details	Displays the log details.

Log types (Global Search)

Global search includes a wide array of log types to help you analyze your log data and identify trends. Use the log type tools to narrow your search, isolate data, or compare two log searches at the same time.

Histogram

The log *Histogram* displays the number of logs collected within the time range selected in the *Global Search* settings. You can use your mouse to select a custom time range and zoom in on the log display



Event Inspector

Click the *Query inspector* button to open the *Query Inspector* panel, and then run the query. The Query Inspector panel shows the plain log search result object. The *Stats* tab shows the query information such as processing time, request time, etc.

The screenshot shows the *Query Inspector* panel. The panel displays the search results for the query 'FortiGate Traffic' in the 'LocalCluster'. The *Query Inspector* tab is active, showing the search result object. The object contains the following fields:

```

Object
  request: Object
    url: "api/datasources/proxy/1/loki/api
D%22%22%7D&start=1616116830000000000"
    hideFromInspector: false
  response: Object
    status: "success"
    data: Object
      resultType: "streams"
      result: Array[1]
        0: Object
          stream: Object
            adomid: "3"

```

Faceted Search (+, -, focus)

Click a log line to view the *Log Details* panel to display the *Log Labels* and *Parsed Fields*.

Common labels: 3 JSON Default LocalCluster fgt/traffic Limit: 1000 (1000 returned)

```

▼ 2021-03-19 14:48:54 itime=2021-03-19 21:48:54 adomid=3 devid=FGVM010000066103 vd=root id=
e=forward level=notice action=accept policyid=5 sessionid=506830 srcip
t=2 service=DNS app=DNS appcat=unscanned srcintfrole=lan dstintfrole=
=5cccc694-fcbb-51e9-4a66-9ca5193e00c6 srcmac=00:14:c2:26:92:82 master:
com tz=-0700

```

Log Labels:

			adomid	3
			source	JSON
			federationName	Default
			serverName	LocalCluster
			logType	fgt/traffic

Parsed Fields:

				action	accept
				adomid	3
				app	DNS
				appcat	unscanned
				devid	FGVM010000066103
				dstcountry	United
				dstepid	101

The following table describes the function of each icon in the Faceted Search:

Icon	Description
	Shows the value statistics of the label or fields.
	Adds the label or field as a condition in query language. For example, app=DNS.
	Excludes the label or field in the query language. Fore example, app!=DNS.
	Displays only this label or field in the log item. For example, app=DNS.

Time Window

Use the time window to select a time range for your log search. Relative time ranges are provided (*Last 5 minutes* to *Last 7 days*). You can also use the *From* and *To* fields to specify a custom time range. Click the *Back*

and *Forward* buttons to move back and forth in the Time Window.

The screenshot shows a 'Time Window' dialog box. At the top, there is a search bar with a clock icon and a time range: '2021-03-19 14:28:25 to 2021-03-19 14:28:29'. To the left of the search bar are 'Split', 'Back', and 'Forward' buttons. Below the search bar, the dialog is divided into two columns: 'Absolute time range' and 'Relative time ranges'. Under 'Absolute time range', there are 'From' and 'To' input fields with the values '2021-03-19 14:28:25' and '2021-03-19 14:28:29' respectively, and a blue 'Apply time range' button. Under 'Relative time ranges', there is a list of options: 'Last 5 minutes', 'Last 15 minutes', 'Last 30 minutes', 'Last 1 hour', 'Last 3 hours', 'Last 6 hours', 'Last 12 hours', 'Last 24 hours', 'Last 2 days', and 'Last 7 days'. At the bottom of the dialog, there is a 'Browser Time' section showing 'United States, PDT' and a 'Change time zone' button set to 'UTC-07:00'. A help message is displayed in a light gray box: 'It looks like you haven't used this timer picker before. As soon as you enter some time intervals, recently used intervals will appear here. Read the documentation to find out more about how to enter custom time ranges.'

Search History

Click *Query history* to show the log search history. After the query is run, you have the option to comment ☐, favorite ☆, clone 📄, or delete 🗑 the query.

The screenshot shows the 'Explore' interface with the following configuration:

- Cluster: LocalCluster
- Traffic: FortiGate Traffic
- Labels: {adomid="3", appcat="Web.Client", logType="fgt/traffic"}
- Line limit: auto
- Refresh: 3.9s

Below the search bar, there are tabs for 'Query history', 'Starred', and 'Settings'. The 'Query history' tab is active, showing a filter history for 'today' and 'a week ago'. Three queries are listed for March 18:

- FAZ-BD: {adomid="3", appcat="Web.Client", logType="fgt/traffic"} (Run query)
- FAZ-BD: {adomid="3", logType="fgt/traffic"} (Run query)
- FAZ-BD: {logType="fgt/traffic", adomid="3"} (Run query)

Split View

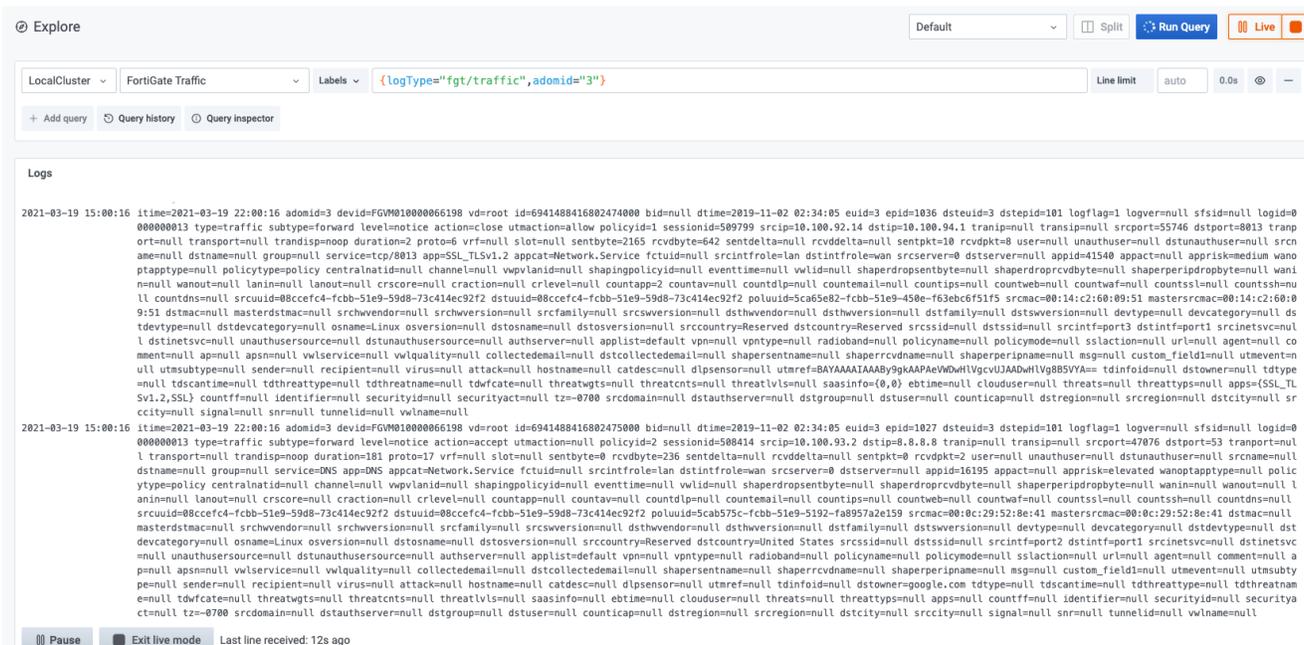
Click the *Split* button to enable *Split View* mode. Split View displays two log search panes so you can search different content and view the results at the same time.

The screenshot shows the 'Explore' interface in Split View mode with two panes:

- Left Pane:**
 - Traffic: FortiGate Traffic
 - Labels: {logType="fgt/traffic", adomid="3"}
 - Line limit: auto
 - Refresh: 3.4s
 - Log chart: Shows a single bar at 14:49:00 with a value of approximately 1000.
 - Log entry: 2021-03-19 14:48:54 itime=2021-03-19 21:48:54 adomid=3 devid=FGW010000066103 vd=root id=694148548763477 6000 dtime=2019-11-02 02:30:19 euid=3 epid=1065 dsteid=3 dstepid=101 logflag=1 log1 d=0000000013 type=traffic subtype=forward level=notice action=accept policyid=5 sess ionid=506830 srcip=10.200.1.17 dstip=8.8.8.8 srcport=46930 dstport=53trandisp=noop duration=181 proto=17 sentbyte=116 rcvbyte=160 sentpkt=2 rcvdpkt=2 service=DNS app= DNS appcat=unscanned srcintfrole=lan dstintfrole=wan srcserver=0 policytype=policy s rcuid=88ccef4-fcb-51e9-59d8-73c414ec92f2 dstuid=88ccef4-fcb-51e9-59d8-73c414ec 92f2 poluid=5cc6694-fcb-51e9-4966-9ca5193e08c6 srcmac=00:14:c2:26:92:82 mastersrc
- Right Pane:**
 - Traffic: FortiGate DNS
 - Labels: {logType="fgt/dns", adomid="3", serverName="LocalCluster"}
 - Line limit: auto
 - Refresh: 2.2s
 - Log chart: Shows two bars at 14:25 and 14:40 with values of approximately 300 and 400 respectively.
 - Log entry: 2021-03-19 14:47:49 itime=2021-03-19 21:47:49 adomid=3 devid=FGW010000066102 vd=root id=694148520846190 2000 dtime=2019-11-02 02:29:57 euid=3 epid=1044 dsteid=3 dstepid=101 logid=15010548 03 type=utn subtype=dns level=warning action=redirect sessionId=508002 policyid=1 sr cip=10.100.92.15 dstip=8.8.8.8 srcport=54449 dstport=53 proto=17 cat=26 xid=50468 qt ypeval=1 srcintfrole=lan dstintfrole=wan ipaddr=(208.91.112.55) srcintf=port3 dstint f=port1 profile=default qname=99.goodyouxi.com qtype=A qlclass=IN catdesc=Malicious W ebsites eventtype=dns-response msg=Domain belongs to a denied category in policy tz=-0700

Live Streaming Search

Click the *Live* button at the top-right corner of your search to enable *Live Log* search. Live Log search displays search results in real-time. Click *Pause* to pause the real-time results, or click *Exit live mode* to return to normal mode.



Cross-Cluster Search Federation

Cross-Cluster search allows you to run searches against one or more remote FortiAnalyzer-BigData clusters. To perform a cross-cluster search, you must have a Search Federation configured. Click the *Federation* menu to open the Federation management UI.

🏠 / Federation
🔔 ¹ ↕

Incoming Federation Request

Allow incoming federation request C

From Server	From User	To User	Status	Received Time	Accepted Time	Actions
10.105.101.59	admin	All Users	Accepted	2/9/2021, 8:12:33 PM	2/9/2021, 8:12:53 PM	Remove

Outgoing Federation Request

[+ New Request](#) C

To Server Name	To Server Address	To User	Status	Sent Time	Confirmed Time	Actions
• 10.105.101.4	10.105.101.4	All Users	Confirmed	2/2/2021, 8:14:04 PM	2/2/2021, 8:14:18 PM	Remove
• 10.105.101.59	10.105.101.59	All Users	Confirmed	2/4/2021, 12:38:27 PM	2/4/2021, 12:38:39 PM	Remove

Search Federation

[+ Add Federation](#) C

Federation Name	Created At	Last Updated	Actions
[-] My Search	2/2/2021, 8:14:33 PM	2/2/2021, 8:14:33 PM	Edit Delete

Server Name	Server Address
• 10.105.101.4	10.105.101.4

Create a new Search Federation (Example)

In the following example, a user on a local cluster (10.106.2.166) wants to create a *Search Federation* using a remote cluster (10.105.101.59).

1. On the remote cluster (10.105.101.59), click the *Allow incoming federation request* button to allow the incoming federation request.

 / Federation

Incoming Federation Request

Allow incoming federation request

2. On the local cluster (10.106.2.166), click the *New Request* button under *Outgoing Federation Request* section and configure the request.

New Federation Request ✕

* Target Server Name:

* Target Server Address:

Auto-confirm Token:

NOTE: This feature only works when sites certificates are trusted

If the remote cluster is using a self-signed certificate, you may see the following dialog for certificate verification. Click *Accept* to send the request.

Verify Certificate X

In order to communicate with other server, the following certificate must be reviewed for correctness, and accepted if deemed valid. Do you wish to accept the certificate as detailed below?

Certificate

Version	3
Serial Number	49 C6 97 7F E4 F6 AE 1F

Subject

Subject Name	EMAILADDRESS=support@fortinet.com, L=Sunnyvale, ST=California, C=US, OU=FortiAnalyzer, O=Fortinet, CN=FBD45FTG19000005
Common Name	FBD45FTG19000005
Organization	Fortinet
Organization Unit	FortiAnalyzer
Locality	Sunnyvale
State	California
Country/Region	US

Issuer

Issuer Name	EMAILADDRESS=support@fortinet.com, L=Sunnyvale, ST=California, C=US, OU=FortiAnalyzer, O=Fortinet, CN=FBD45FTG19000005
-	..
-	-----

After the request is sent, the you will see a *Pending* item created in the table. You can cancel this item any time.

Outgoing Federation Request

+ New Request

To Server Name	To Server Address	To User	Status	Sent Time	Confirmed Time	Actions
10.105.101.4	10.105.101.4	All Users	Confirmed	2/2/2021, 8:14:04 PM	2/2/2021, 8:14:18 PM	Remove
10.105.101.59	10.105.101.59	All Users	Pending	3/22/2021, 4:42:45 PM	N/A	Cancel

- Go back to *Federation Management* in the remote cluster (10.105.101.59), and click *Accept* to accept this federation request, or click *Ignore* to ignore the request.

Incoming Federation Request

Allow incoming federation request (expired after 12m)

From Server	From User	To User	Status	Received Time	Accepted Time	Actions
10.106.2.166	admin	All Users	Pending	3/22/2021, 4:42:55 PM	N/A	Accept Ignore

- On the local cluster (10.106.2.166), click the *Add Federation* button of the *Search Federation* section, and select the *Federation Servers*.

Add Search Federation ✕

i Add Search Federation to allow cross-cluster search

* Federation Name:

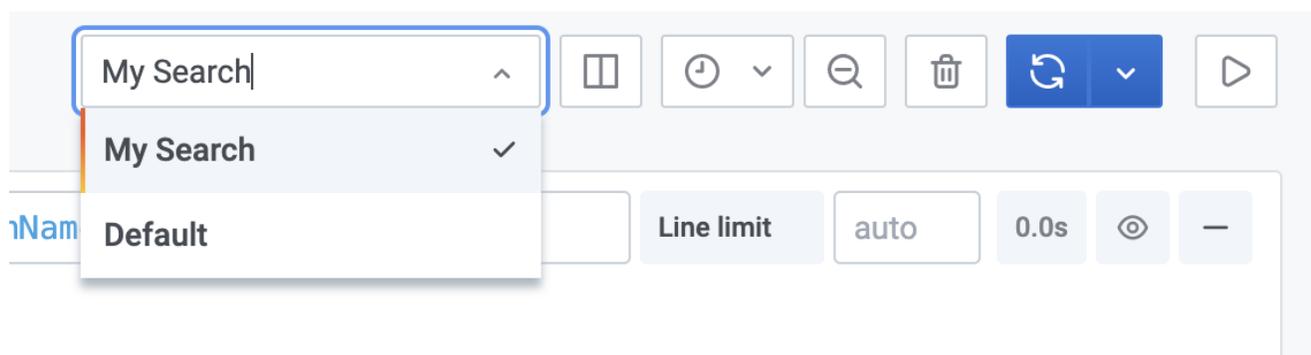
Federation Servers:

- 10.105.101.4
- 10.105.101.59

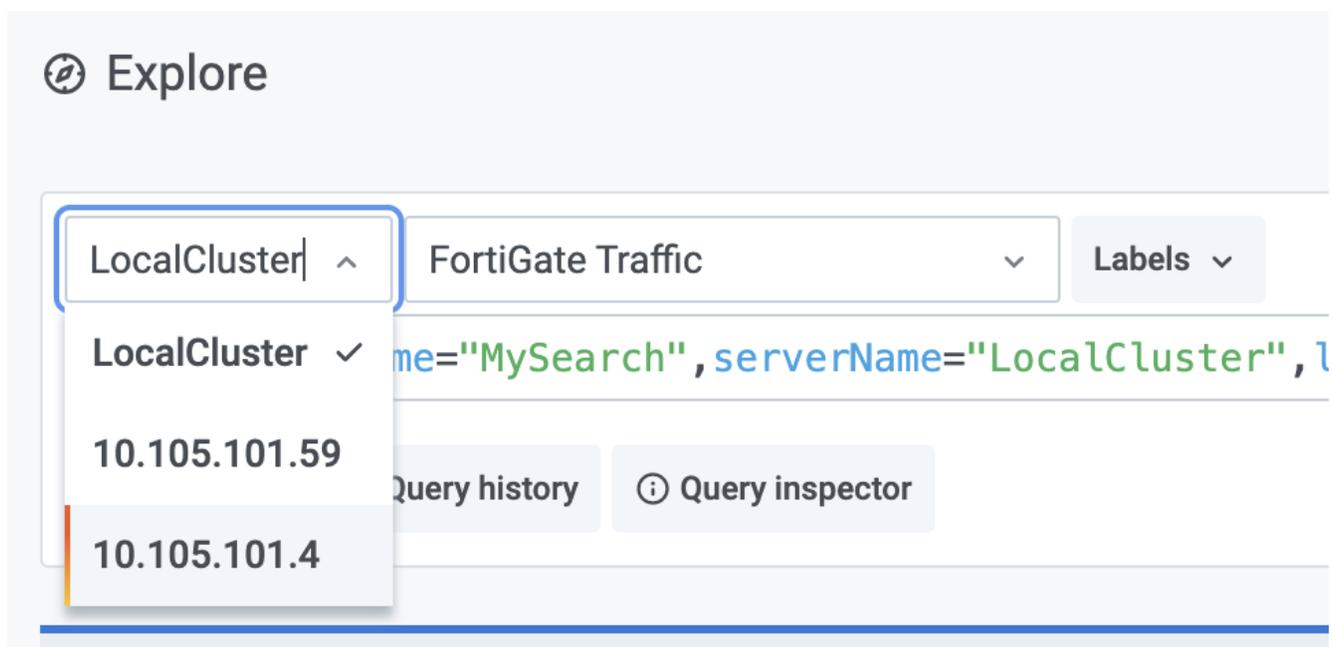
Search Federation			
Federation Name	Created At	Last Updated	Actions
<input type="checkbox"/> My Search	2/2/2021, 8:14:33 PM	2/2/2021, 8:14:33 PM	Edit Delete
Server Name		Server Address	
• 10.105.101.4		10.105.101.4	
• 10.105.101.59		10.105.101.59	

Search with Federation

To search with a different cluster, select a *Search Federation* that contains multiple servers (for example, *My Search*), and then select the cluster name (10.105.101.59) from the *Cluster* dropdown.



And then select the cluster name (e.g. 10.105.101.59) in the *Cluster* dropdown.



You can also add another query for a remote cluster by clicking the *Add query* button. The search returns the results from both servers.



Log Query Language (LogQL)

To create custom queries, use *LogQL* in the log query input box of the *Global Search*. LogQL can be considered a distributed *grep* that aggregates log sources. LogQL uses labels and operators for filtering.

A basic log query consists of two parts:

- Log stream selector
- Log pipeline

Log Stream Selector

The log stream selector determines which log streams should be included in your query results. The stream selector is comprised of one or more key-value pairs, where each key is a log label and each value is that label's value. The log stream selector is written by wrapping the key-value pairs in a pair of curly braces:

```
{logType="fgt/traffic", adomid="3"}
```

In the example above, all log streams that have a label of `logType`, whose value is `fgt/traffic` and a label of `adomid` whose value is 3 will be included in the query results. This will match any log stream whose labels contains at least 3 for their `adomid` label. If there are multiple streams that contain that label, logs from all of the matching streams will appear in the results.

The `=` operator after the label name is a label matching operator. The following label matching operators are supported:

- `=`: Equals exactly
- `!=`: Not equal

Log Pipeline

Optionally, the log stream selector can be followed by a log pipeline. A log pipeline is a set of stage expressions chained together and applied to the selected log streams.

A log pipeline can be appended to a log stream selector to further process and filter log streams. This usually consists of one or multiple expressions, each expression is executed in sequence for each log line. If an expression filters out a log line, the pipeline will stop at this point and start processing the next line. An expression is a SQL-where-clause-like condition.

```
{logType="fgt/traffic", appcat="Collaboration"} | osname ILIKE 'windows%' AND dstinetsvc IREGEXP '^.*gmail.*' AND sentbyte > 10000000
```

In the example above, the condition will filter out the FortiGate traffic Collaboration app category log messages when the OS name contains "windows" in the beginning and destination internet service matches Gmail and sent bytes is greater than 10000000(10MB).

The following operators are supported in the Log Pipeline expression:

Operator	Description
<code>=, !=, <, <=, >, >=</code>	Comparison operators.
AND, OR, NOT	Logical operators.
BETWEEN ... AND ...	Compares to both a lower (<code>>=</code>) and upper (<code><=</code>) bound.
IN	Compares an argument value to a set of values and returns TRUE if the argument matches any value in the set. NOT IN reverses the comparison.
LIKE	Comparison operator for STRING , with basic wildcard capability using <code>_</code> to match a single character and <code>%</code> to match multiple characters.
ILIKE	Case insensitive LIKE .
REGEXP	Tests whether an argument value matches a regular expression. Uses the POSIX regular expression syntax where <code>^</code> and <code>\$</code> match the beginning and end of the string: <ul style="list-style-type: none"> • <code>.</code> represents any single character,

Operator	Description
	<ul style="list-style-type: none">• * represents a sequence of zero or more items,• + represents a sequence of one or more items,• ? produces a non-greedy match, and so on.
IREGEXP	Case insensitive REGEX .

Job management and automation

The Jobs page contains a table that displays all jobs in the system, including built-in jobs and custom jobs.

The screenshot shows the 'Jobs' page in FortiAnalyzer. At the top, there are navigation buttons: '+ Create Custom Job', 'Import Job', and 'Export Jobs'. Below these is a table with the following columns: Summary, Job Type, Schedule, Last Job Status, Last Result, Last Job Updated, Create Time, and Actions. The table lists several built-in jobs:

Summary	Job Type	Schedule	Last Job Status	Last Result	Last Job Updated	Create Time	Actions
Storage Group Backup	Build-in	Manual	Failed		4/6/2020, 8:45:52 PM	4/6/2020, 8:45:44 PM	Run
Data Appendix	Build-in	0 0 0/4 ? ***	Success		4/22/2020, 12:00:45 PM	4/1/2020, 10:31:04 PM	Run Configure
Facet Formation - Reports	Build-in	0 10/30 * ? ***	Success		4/22/2020, 12:23:44 PM	4/1/2020, 10:31:04 PM	Run Configure
Facet Formation - FortiView	Build-in	0 0/5 * ? ***	Success		4/22/2020, 12:23:00 PM	4/1/2020, 10:31:04 PM	Run Configure
Data Retention	Build-in	0 30 * ? ***	Success		4/22/2020, 11:30:28 AM	4/1/2020, 10:31:04 PM	Run Configure
Data Rebalance	Build-in	0 0 0 ? * TUE,THU,SAT *	Success		4/21/2020, 12:14:23 AM	4/1/2020, 10:27:14 PM	Run Configure

The Jobs table contains the following columns:

Column header	Description
Summary	The name or short description of a job. You can click the summary to view its execution history (see Job history on page 68).
Job Type	There are two types of jobs: <ul style="list-style-type: none"> • Built-in: Pre-configured system jobs. • Custom: Job created by an administrator.
Schedule	Shows how often the job is run.
Last Job Status	Indicates the status of the job: <ul style="list-style-type: none"> • Success: The job execution successful. • Failed: The job execution failed. • Running: The job is currently executing. • Queued: The job has been put into an execution queue and will be executed shortly. • Timeout: The job execution has timed out. • Aborted: The job execution has been interrupted. This status usually occurs when the user manually aborts. • Skipped: The job has been skipped. This status usually occurs when a previously executed job is still running and its job configuration does not allow concurrent jobs.
Last Job Result	View the last job execution result by clicking Job Result

Column header	Description
Last Job Updated	When the job was last run.
Create Time	When the job was first created.
Actions	You can perform two actions on the health test: <ul style="list-style-type: none"> • Run: Manually launch a job execution. • Configure: Change a job's configurations. • Delete: Delete a job and the job's history.

Job history

To access the Job History page and see the job execution records, click its Job Summary link.

🏠 / Jobs / Data Retention 🔔 ⚙️

Run Job 🔍 ☰

Summary	Status	Fired Time	Triggered By	Duration	Result	Actions
#4/22/2020, 11:30:00 AM	🟢 Success	4/22/2020, 11:30:00 AM	System	28.8s	📄	View Config Delete
#4/22/2020, 10:30:00 AM	🟢 Success	4/22/2020, 10:30:00 AM	System	28.9s	📄	View Config Delete
#4/22/2020, 9:30:00 AM	🟢 Success	4/22/2020, 9:30:00 AM	System	29.1s	📄	View Config Delete
#4/22/2020, 8:30:00 AM	🟢 Success	4/22/2020, 8:30:00 AM	System	29.0s	📄	View Config Delete
#4/22/2020, 7:30:00 AM	🟢 Success	4/22/2020, 7:30:00 AM	System	28.9s	📄	View Config Delete
#4/22/2020, 6:30:00 AM	🟢 Success	4/22/2020, 6:30:00 AM	System	29.3s	📄	View Config Delete

< 1 2 3 4 5 ... 10 >

You can view records of the job's execution result, job configurations, or even delete the record.



FortiAnalyzer-BigData only saves the last 300 records for job execution results

Built-in automation jobs

FortiAnalyzer-BigData has the following default built-in jobs:

Built-in jobs	Description
Data Appendix	Created automatically when storage pools are created. This job generates the list of available sub-types of FortiGate Event logs for LogView.
Data Maintenance	This job computes the data statistics for metadata tables to improve query performance.
Data Retention	Created automatically when storage pools are created. This job is used to apply data retention policies for the Storage Pool which marks the old data for deletion and makes space for future data.
Facet Diagnosis	This job diagnoses Facet Formation processes and remedies any detected errors.
Facet Formation - Report	Created automatically when storage pools are created. This job generates the pre-aggregated facets to speed up Report queries.
Facet Formation - FortiView	Created automatically when storage pools are created. This job generates the pre-aggregated facets to speed up FortiView queries.
Metadata Backup	This job periodically backs up metadata.
Storage Pool Backup	Created automatically when Storage Pool backup is scheduled (see Data backup on page 79).

Custom automation jobs

You can create or import custom jobs by using built-in or custom templates rendered as an Ansible playbook.

To create a custom automation job:

1. In the top-left corner of the Jobs page, click *Create Custom Job*.
The Create Custom Job dialog box loads.

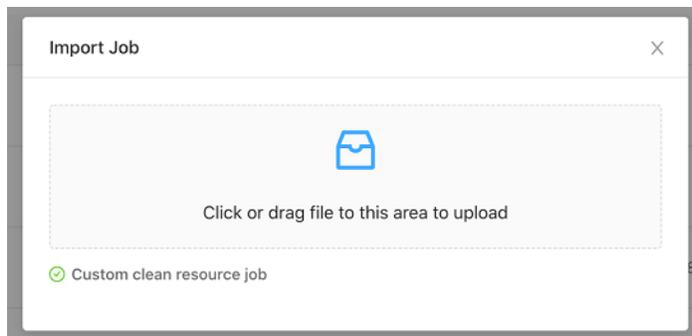
2. Complete the following fields:

Field name	Description
Job Summary	Enter the job description.
Template	Select a job template. For templates that have additional fields to fill out, see Custom job templates on page 71 .
Schedule	Select a scheduling timer: <ul style="list-style-type: none"> • Manual: The job will not be executed until you manually launch it. • Daily: The job is scheduled to run on a daily basis. Select a run time and enable the Enable Job toggle so the schedule takes effect. To pause the job schedule, disable the toggle. • Advanced: Supports standard cron expressions. You can use predefined cron expressions to schedule a run every 30 minutes, every hour, every 12 hours, and more. Switch the Enable Job toggle to enable so the schedule takes effect. To pause the job scheduling, disable the toggle.
Allow Concurrent Job	Enable to allow multiple jobs to run at the same time.
Enable Timeout	Enable to define job timeout.

3. When you finished configuring your job, click *Create*.

To import custom jobs:

1. In the top-left corner of the Jobs page, click *Import Job*. The Import Job dialog box loads.



2. Drag or select the file you want to import into the dialog box.

To export multiple custom jobs:

1. From the Jobs page, select the jobs you want to export.
2. In the top-left corner of the Jobs page, click *Export Job*. The Confirm Export Job dialog box loads.
3. Click *Confirm* to export your jobs.

Custom job templates

When you select a template for your custom job, you might need to fill out additional fields depending on the template you select. The following templates require additional configuration before you can apply them.

Backup Table Validation

The Backup Table Validation template is used to verify the data integrity of the backup data at the selected location.

* Template:

* Storage Group:

* HDFS Url:

Select the storage pool and enter the Hadoop Distributed File System (HDFS) URL for the backup location.

Custom Template

Custom templates are used to create the content for custom jobs for when built-in jobs don't meet your specific needs. You can create custom templates to operate the host, collect information, take actions, and more.

Custom templates require you to use the Ansible playbook YAML format to define the content. For information about Ansible specifications, refer to the [official Ansible documentation](#).

The following example template collects the disk usage of the BigData Controller and sends it to a Slack channel:

```
- name: Collect disk usage and send to slack
  hosts: controllerIp
  vars:
    - slack_url: "https://hooks.slack.com/services/xxxxxxx" # your slack app webhook url
  tasks:
- name: Collect disk usage
  command:"df -h"
  register: result
- name: Send to slack
  uri:
    url:"{{ slack_url }}"
    body:'{"text": "{{ result.stdout }}"}'
    body_format: json
    method: POST
```

The follow table shows all the Ansible inventory group names you can use as hosts values in your playbook and template. Those values are pre-populated in the Ansible inventory and are automatically applied with each execution.

<ul style="list-style-type: none"> • hdfs_datanode • hdfs_namenode • kudu_tserver • kudu hive_metastore • zookeeper • kafka_broker • impala_catalog • impala • impala_statestore • yarn_nodemanager • yarn_resource_manager • spark_history_server 	<p>These inventory groups can be used to select the host(s) that have the named services running.</p> <p>For example, using “host: kudu_tserver” in your playbook allows it to be executed on all hosts has kudu-tserver instance.</p>
<ul style="list-style-type: none"> • hdfs_datanode_reachable • hdfs_namenode_reachable • kudu_tserver_reachable • kudu_reachable • hive_metastore_reachable • zookeeper_reachable • kafka_broker_reachable • impala_catalog_reachable • impala_reachable • impala_statestore_reachable • yarn_nodemanager_reachable 	<p>These groups can be used to select one of the reachable hosts that belong to the named service.</p> <p>For example: kudu has instances spreading on 3 hosts, and “hosts:kudu_reachable” will randomly return one that is reachable at the execution time.</p>

<ul style="list-style-type: none"> reachable yarn_resource_manager_reachable spark_history_server_reachable 	
<ul style="list-style-type: none"> metastore datanode master 	These groups can be used to select hosts the belong to the named role.
<ul style="list-style-type: none"> metastore_reachable datanode_reachable master_reachable 	These groups can be used to select a random host that is reachable at the execution time, from the ones with the named role.
<ul style="list-style-type: none"> controllerlp 	This group can be used to the BigData Controller host.

In addition to these groups, you can also use the host name shown in the Hosts page to directly select a particular host for the playbook execution.

Data Log Type Appendix

The Data Log Type Appendix is run to re-generate the list of available log types for LogView.



This is a resource intensive operation. Run this only if the available log types sidebar of LogView is not working properly.

Docker System Prune

The Docker System Prune template is run to remove all unused docker containers, networks, and images (both dangling and unreferenced) to clear disk space.

Facet Formation Manual Run

The Facet Formation Manual Run enables you to manually run a facet formation. Run this job only when the FortiView query performance is exceptionally slow.

* Template:

* Storage Group:

* Mode: From Beginning Custom Time

Time: Hour(s)

First, select a storage pool, and then select the time to do facet formation. You can choose between starting the facet formation from the beginning, or from a specific time.

HDFS Safemode Leave

The HDFS Safemode Leave template enables you to leave the HDFS safe mode from an unexpected shutdown.

Hive Metastore Backup

The Hive Metastore Backup template creates a backup of the data in Hive Metastore and saves it to an HDFS location.

Hive Metastore Restore

The Hive Metastore Restore template restores the data in Hive Metastore from an HDFS location.

Kafka Deep Clean

The Kafka Deep Clean template deep cleans Kafka topics and reinstalls Kafka (see [How to recover from an unhealthy service status on page 122](#)).

Kafka Rebalance

The Kafka Rebalance template rebalances the data load across the hosts. This is useful for when a Kafka node is decommissioned or when a new Kafka node joins or leaves the cluster. It includes replica leadership rebalance and partition rebalance. For more information, see [Scaling FortiAnalyzer-BigData on page 103](#).

NTP Sync

The NTP Sync template performs a manual NTP time sync on all the BigData hosts. Run this job when Kudu time synchronization is unsynced (see [How to recover from an unhealthy service status on page 122](#)).

Purge Data Pipeline

This job resets the watermark and performs a clean restart of the pipeline.



Any unprocessed data will be lost (see [How to recover from an unhealthy service status on page 122](#)).

Data management

FortiAnalyzer-BigData uses Storage Pools to manage disk space. The *Data* page contains the Storage Pools table where you can monitor the group's status, manage storage policies and jobs, and backup or restore data. A default *Root* storage pool is included with FortiAnalyzer-BigData.

Storage Pool	Storage Pool Type	Status	ADOMs	Age	Data Usage	Actions
Root	Default	Ready	16	0 Day / 60 Days	320.5 GB / 162.2 TB	Manage Data Policy More

ADOM Name	Status	Device Type	Description
root	Ready	Fabric	
FortiCarrier	Ready	FortiGate	
FortiMail	Ready	FortiMail	
FortiAnalyzer	Ready	FortiAnalyzer	
FortiWeb	Ready	FortiWeb	
FortiCache	Ready	FortiCache	
FortiClient	Ready	FortiClient	
Syslog	Ready	Syslog	
FortiManager	Ready	FortiManager	

The Storage Pool table contains the following columns:

Column header	Description
Storage Pool	The name of storage pool. You can expand each storage pool to display all the ADOMs in that group.
Storage Pool Type	The type of storage pool, which indicates the storage strategy.
Status	Indicates the status of the storage pool. <ul style="list-style-type: none">• Ready: The storage pool is ready for use.• In Progress: The storage pool is being created and is not yet ready for use.• Failed: The storage pool creation failed.
ADOMs	The number of ADOMs in that storage pool.
Age	The age of the storage pool and the number of days the logs will be stored.
Data Usage	The amount of data in use.
Actions	You can perform the following actions on a storage pool: <ul style="list-style-type: none">• Manage Storage Policy: Determine how long and the maximum size you want to store the data, and when to do a data rollover. For more information, see Manage storage policy on page 78.• Manage Job: Manage jobs in that storage pool.• Backup: Create a backup of that storage pool.• Manage Backups: Show and delete the backup data.

Column header	Description
	<ul style="list-style-type: none"> • Restore: Restore data. • Migrate ADOM: Migrate ADOM from one storage pool to another. • Delete Storage Pool: Delete the storage pool. This action is only available if no ADOM exists in the storage pool.

To create a Storage Pool:

1. Click *Create Storage Pool*.
2. Configure the storage pool settings.

Option	Description
Storage Pool Name	Enter a name for the storage pool.
Description	Enter a description of the storage pool.
Storage Pool Type	Select the type of the storage pool. <ul style="list-style-type: none"> • Default: Default strategy optimized for common use cases when data are distributed across ADOMs. • Skewed: Optimized when the majority of the data are distributed in only one or a few ADOMs.
Keep Logs For	Enter the number days the logs are to be stored.
Allocated	Enter the amount of memory allocated for storing the logs.

New Storage Pool X

* Storage Pool Name:

Description:

* Storage Pool Type:

i Default strategy optimized for common use cases when data are distributed across ADOMs

Keep Logs For:

Allocated:

Maximum Available: 21.8TB

3. Click *Create*.

To assign a Storage Pool to an ADOM:

1. Create a new ADOM.
2. From the *Storage Pool* dropdown, select a storage pool or click the Plus (+) sign to create a new pool.

Create ADOM

Name: AdomTest

Type: FortiAnalyzer

Storage Pool: Storage 1

Description: [Search icon]

Devices

Name	IP Address	Platform
No record found.		

Cache Utilization

Allocated: 51200 MB Maximum Available: 1.5 TB

Analytics: Archive: 50% 50% Modify

Alert and Delete When Usage Reaches: 90%

OK Cancel

Manage storage policy

You can manage the storage policy of each storage pool from the *Actions* column on the Data page.

To manage a data policy:

1. From the Data page, select a Storage Pool and click *Actions > Manage Data Policy*. The *Manage Data Policy* dialog opens.

Manage Data Policy

Data Policy

Keep Logs For: 60 day(s)

Cancel OK

2. In the *Keep Logs For* field, select the number of days you want to store the data in the system. FortiAnalyzer-BigData removes the data from the system after the selected number of days.
3. Click *OK*.

To manage disk allocation:

1. From the Data page, click *Manage Disk Allocation*.
The *Manage Disk Allocation* dialog opens.

Storage Pool	Disk Allocated
Root	0.1TB - 2TB - 3.9TB 0TB
S1	0.1TB - 1TB - 3.9TB 0TB
S2	0.1TB - 0.9TB - 3.9TB 0TB
Total	3.9 TB / 3.9 TB

2. For each storage pool, move the slider in the *Disk Allocated* column to configure the maximum amount of data you want to store in the system.
FortiAnalyzer-BigData removes the data from the system after the data size exceeds the selected number.
3. Click *OK*.

Data backup

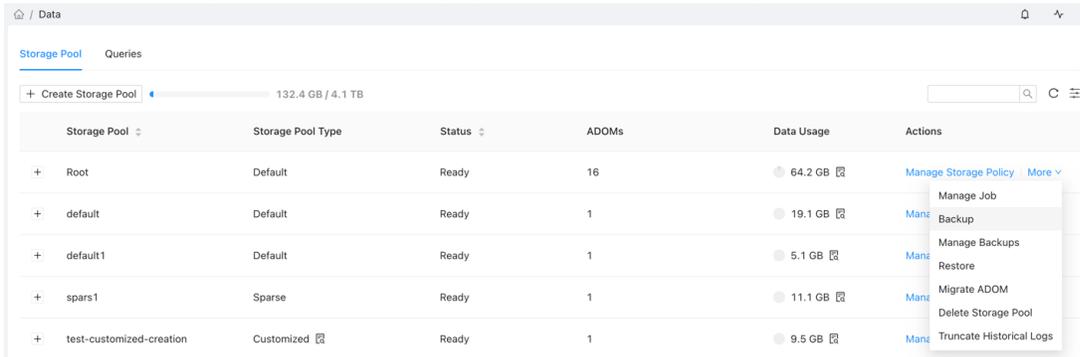


The cluster-level high availability semantics guarantee that data stored in a Storage Pool can tolerate one blade lost at a time (data are replicated three times and distributed across the cluster hosts). We recommend backing up data to an external HDFS cluster for additional data-center-level redundancy.

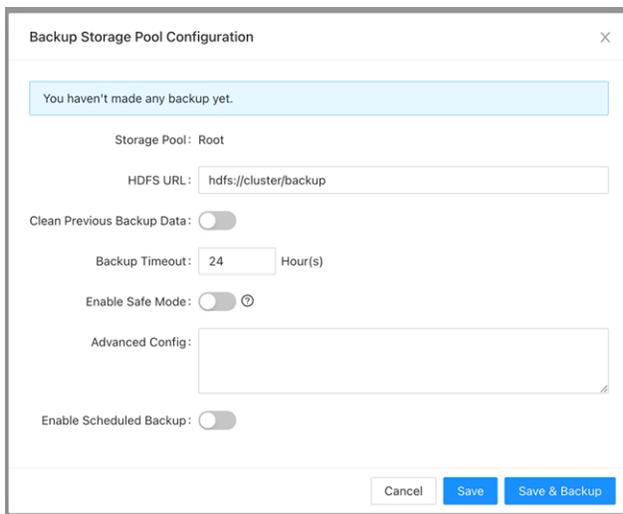
FortiAnalyzer-BigData supports disaster recovery and data portability. You can back up all the data within a Storage Pool to Hadoop Distributed File System (HDFS) in Parquet file format.

To back up data:

1. From the Data page, locate the storage pool you want to back up and select *Actions > Backup*.



The Backup Storage Pool Configuration dialog loads with the following fields:



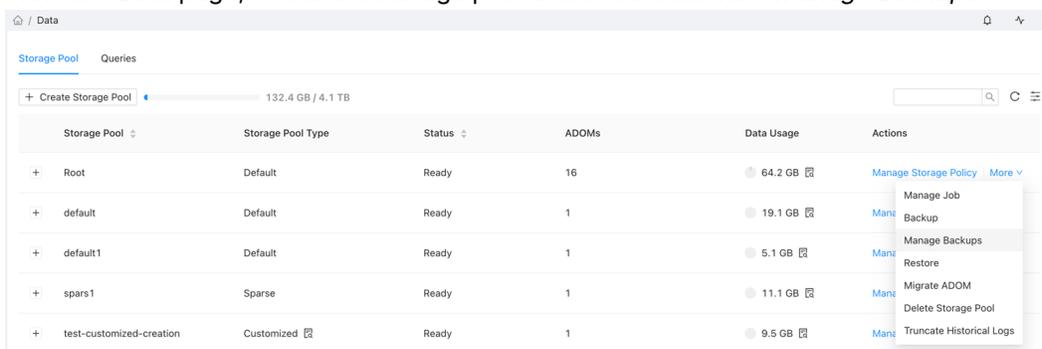
Field name	Description
HDFS Url	<p>Defines the target directory of the HDFS cluster. By default, the field is set to the built-in HDFS in the Security Event Manager. We recommend backup to an external HDFS.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  <p>If the URL is configured to an external HDFS cluster, all its hosts must be made accessible by the FortiAnalyzer-BigData hosts (see Backup and restore to external HDFS on page 97).</p> </div>
Clean Previous Backup Data	<p>Enable to delete any previous backup data and start a new backup. Do not enable if you want to create an incremental backup.</p>
Backup Timeout	<p>Enter the number of hours before the backup job times out. After the timeout, the job will abort.</p>

Field name	Description
Enable Safe Mode	By default, the normal backup job processes multiple tables in parallel and ignore any intermediate errors. Enable Safe Mode to back up the Storage Pool tables sequentially and to fail early if any error occurs.
	 <p>This mode may take longer to complete the back up, so only enable Safe Mode when the normal backup job fails.</p>
Advanced Config	These configurations define the resources used for the job. Normal users should keep the default configurations.
Enable Scheduled Backup	Enable so the backup can be scheduled automatically.

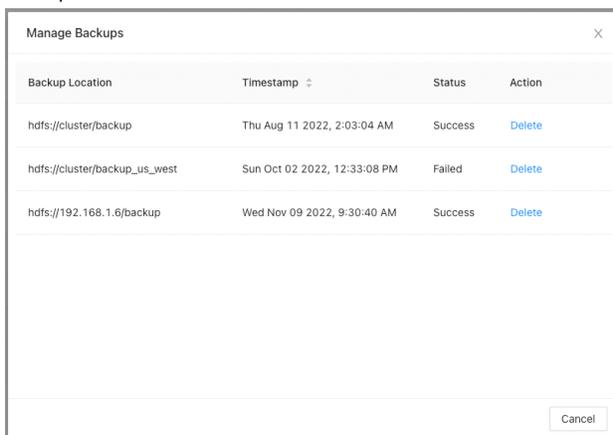
- When you are finished, click *Save & Backup* to begin the backup process.
- You can monitor the status of your backup by navigating to *Jobs > Storage Pool Backup*.

To delete backup data:

- From the Data page, locate the storage pool and select *Actions > Manage Backups*.



- The *Manage Backups* dialog displays the backup location, timestamp, and status. To delete the backup data, select *Action > Delete*.



Incremental backups

We recommend that you create incremental backups by consistently backing up new data to the same HDFS directory.

The first time a backup job is run, a full backup of the storage pool data will be saved to the HDFS directory. Subsequent runs will perform incremental backups which only contain the rows that have changed since the initial full backup. Thus, the subsequent backups will be faster and more efficient.

To create manual incremental backups:

If you have already created a previous backup, you can manually create an incremental backup against it.

1. From the navigation bar, go *Jobs* and click *Storage Pool Backup* to view all the completed backups.
2. Select the backup which you want to create an incremental backup against and click *View Config*.

The Job Instance Configuration dialog loads with the following fields:

Job Instance Configuration		X
Template :	Storage Pool Backup	
Storage Pool	Root	
HDFS Url :	hdfs://cluster/backup	
Backup Timeout :	24	
Clean Previous Backup Data :	false	
Enable Safe Mode :	false	
Advanced Config :		

3. In the *HDFS Url* field, copy the URL.
For example: hdfs://cluster/backup/7o7T
4. Go to *Data* and select the same Storage Pool as the previous backup, and click *Actions > Backup*.
5. In the *HDFS URL* field, paste in the HDFS Url copied from step 3.



You can check the number of existing backups in the Backup Storage Pool Configuration dialog.

Backup Storage Pool Configuration X

You have made 1 backups, Click [here](#) to check details.

Storage Pool: Root

HDFS URL:

Clean Previous Backup Data:

Backup Timeout: Hour(s)

Enable Safe Mode: ⓘ

Advanced Config:

Enable Scheduled Backup:

6. Ensure the *Clean Previous Backup Data* option is disabled so you do not clean any previous backup data, allowing this backup to be incremental.



You can enable this option to make a full backup to the HDFS directory, however, a full backup job will be more time consuming than an incremental backup.

7. When you are finished, click *Save & Backup* to begin the backup process.

To create scheduled incremental backups:

You can also schedule incremental backup jobs by enabling the *Enable Scheduled Backup* option. This schedules incremental backup jobs to the HDFS you set. Fortinet strongly recommends scheduling maintenance jobs at off-peak hours.

Data restore



Restoring data requires you to drop all tables in the storage pool. Be cautious when selecting your configurations.

To restore data from a backup

1. From the navigation bar, go to Data and select the Storage Pool you want to restore data for.

- In the Storage Pool row, click *Actions > Restore*.
The Restore Storage Pool Configuration dialog loads.

Restore Storage Pool Configuration X

Storage Pool: Root

* Select Backup: v

* HDFS URL:

Backup Tables:

Backup Timestamp: 📅

Restore Timeout: Hour(s)

Enable Safe Mode: ?

Advanced Config:

Restore storage pool action will pause the log input (Ingestion and Pipeline will be stopped), and **DELETE** all data (**CANNOT** be undone) in current storage pool. Please type **restore** to confirm:

- Configure the following fields:

Field name	Description
Select Backup	Select the backup type you used. If the data is from an external system, select Custom.

Field name	Description
HDFS URL	<p>Defines the location of the backup.</p> <p>If the URL is configured to an external HDFS cluster, all of its hosts must be accessible by of the hosts of the Security Event Manager.</p>
Backup Timestamp	<p>This config can be used to limit the data that you want to restore. It only applies for multi-backups (multiple incremental backups).</p> <p>The following figure shows how multi-backups are restored:</p>
Enable Safe Mode	<p>By default, the normal backup job processes multiple tables in parallel and ignore any intermediate errors. Enable Safe Mode to back up the Storage Pool tables sequentially and to fail early if any error occurs.</p> <hr/> <div style="display: flex; align-items: center;"> <p>This mode may take longer to complete the back up, so only enable Safe Mode when the normal backup job fails.</p> </div> <hr/>
Restore Timeout	<p>Enter the number of hours before the restore job times out. After the hours elapse, the job will abort.</p>
Advanced Config	<p>These configurations define the resources used for the job. Normal users should keep the default configurations.</p>

4. When you are finished, enter restore into the confirmation box to confirm.
5. Click *Restore* to begin the data restoration process.

Migrate an ADOM to another Storage Pool

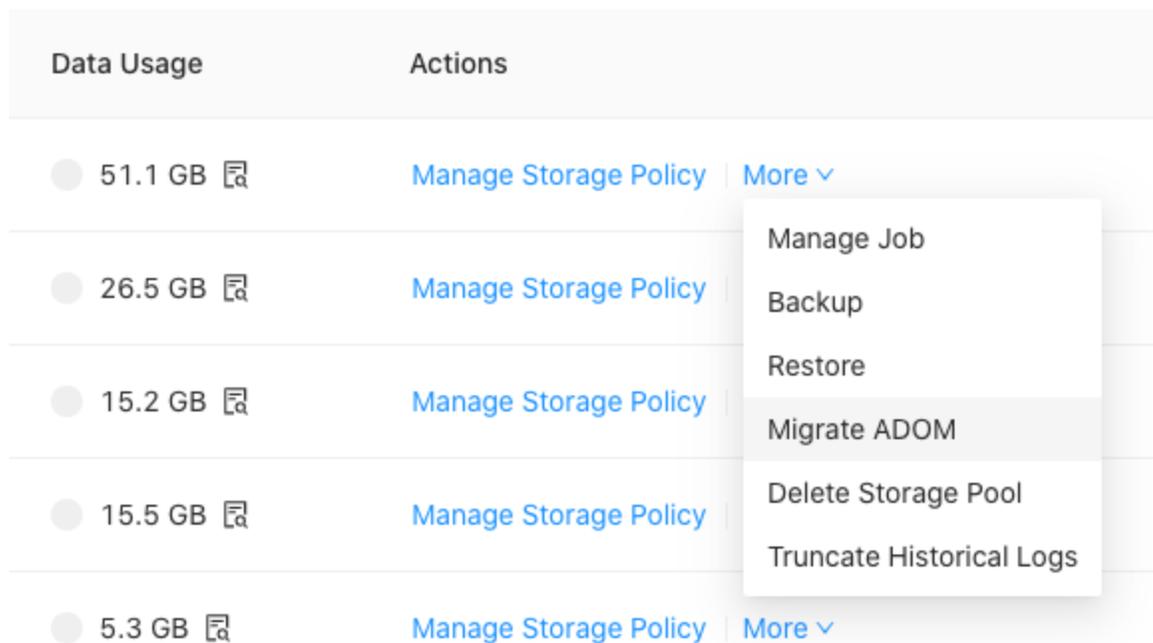
Grouping ADOMs with similar log rates and data retention requirements into a storage pool is recommended. For example, group small ADOMs (in terms of log rate and data volume) into one Storage Pool and larger ADOMs in another. When different sized ADOMs are grouped into one storage pool, the query performance on the smaller ADOMs will be affected by the larger ADOMS.

You may need to move an ADOM to another Storage Pool if its data characteristics change dramatically (e.g. increase in log rate) or data retention requirement changes.

After migration, new incoming logs are sent to the new Storage Pool. Historical logs will be kept in the old Storage Pool and managed by the old data retention policy. Logs can be queried seamlessly even though the data resides in multiple Storage Pools.

To migrate an ADOM to another Storage Pool:

1. Go to *System Settings > Storage Pool*, and locate the Storage Pool that contains the ADOMs to migrate.
2. In the *Actions* column, click *More > Migrate ADOM*. The migration dialog opens.



Data Usage	Actions
● 51.1 GB 	Manage Storage Policy More 
● 26.5 GB 	Manage Storage Policy
● 15.2 GB 	Manage Storage Policy
● 15.5 GB 	Manage Storage Policy
● 5.3 GB 	Manage Storage Policy More 

3. From the *ADOM* dropdown, select the ADOM you want to migrate, and from the *Migrate to* dropdown, select the target Storage Pool.

ADOM Migration

* ADOM: FortiManager

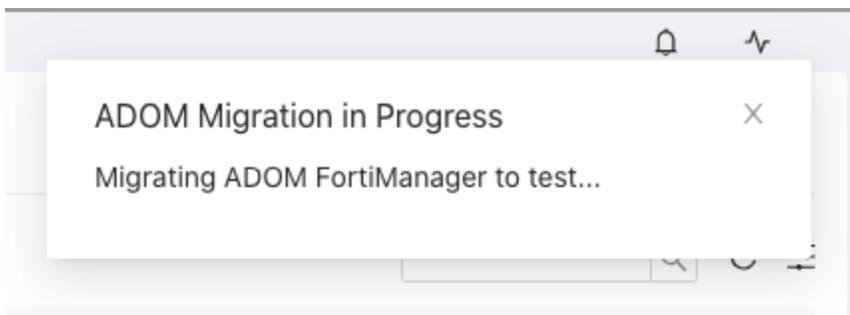
* Migrate to: test

NOTE: The migration will be effective on new incoming logs. Historical logs will be kept in the old Storage Pool. This action will not affect Log View search, but it may affect FortiView query performance temporarily as the system will have to rebuild the cache.

Migrate ADOM **FortiManager** from Storage Pool **Root** to **test**. The migration may take several minutes (will cause no interruption to log ingestion, FortiView or Log View, etc.).

Cancel Migrate

4. Click *Migrate*. The migration progress notification is displayed.



The s ADOM Status changes to *Migrating* in both Storage Pools.

–	Root	Ready
	ADOM Name ▾	Status ▾
	FortiManager	<div style="width: 50%; background-color: #ccc;"></div> Migrating...

–	test	Ready
	ADOM Name ▾	Status ▾
	FortiCache	Ready
	FortiManager	<div style="width: 20%; background-color: #007bff;"></div> Migrating...
	FortiFirewall ⌚	Ready
	FortiProxy	Ready
	FortiFirewall	Ready

When the migration is complete, two ADOM records are created in both Storage Pools. The pool with the clock icon ⌚ indicates the ADOM with historical log data. The pool without the clock icon is the ADOM where the new incoming logs are sent.

To view the time range of historical logs, hover over the clock icon.

	Storage Pool ▾	Status ▾
–	Root	Ready
	ADOM Name ▾	Status ▾
	Historical: 3/15/2022 23:35:31 - 3/16/2022 23:48:1 UTC	Ready
	FortiManager ⌚	Ready

Truncate historical logs of the migrated ADOM

You can truncate the historical logs of an ADOM in a Storage Pool. Once truncated, the historical logs of the ADOM will be marked for deletion. After the logs are truncated the log data is not searchable, however the space is not cleared immediately. The log data will be retained by the data retention policy of the Storage Pool.

To truncate historical logs:

1. Go to *System Settings > Storage Pool*, and locate the Storage Pool that contains the historical ADOM to truncate.
2. In the *Actions* column, click *More > Truncate Historical Logs*. The migration dialog opens.
3. Click *Truncate*.

Truncate Historical Logs of the ADOM [X]

* ADOM: FortiManager [v]

NOTE: The historical logs of the ADOM in this Storage Pool will be marked for deletion and are maintained per the data retention policy. This action will not affect the logs of the ADOM in other Storage Pools. [X]

Cancel Truncate

Bootloader

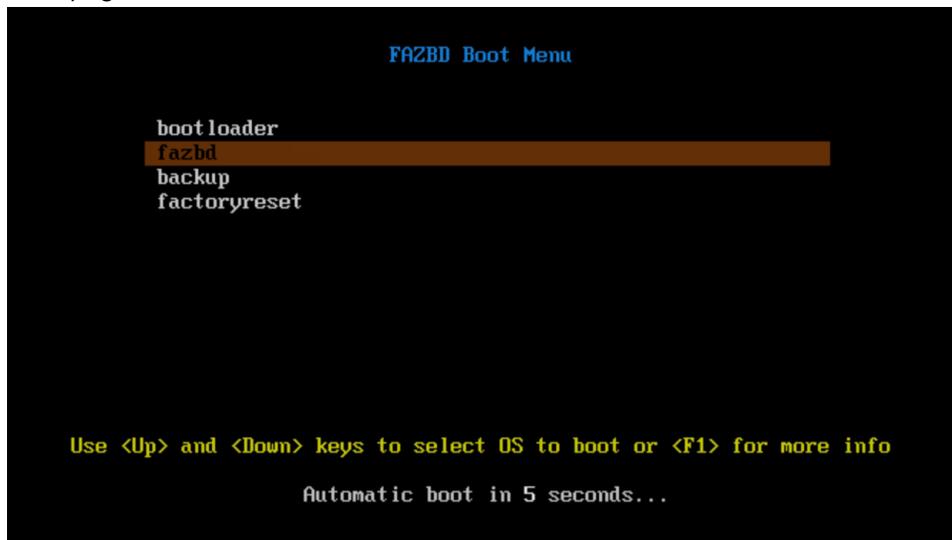
The FortiAnalyzer-BigData Bootloader is a system software that manages the FortiAnalyzer-BigData host's firmware. The Bootloader can be accessed during host reboot. The Bootloader can be accessed on all the hosts.



Improper selection of options in FortiAnalyzer-BigData Bootloader can have an adverse impact on the whole system, and even lead to system failure. Approach these options with great care and when in doubt, err on the side of caution.

To access the Bootloader for a FortiAnalyzer-BigData:

1. Connect to the CMM web management utility.
See *Connect to the Chassis Management Module* in the FortiAnalyzer-BigData Getting Started Guide for your appliance: [4500F](#) or [4500G](#).
2. Select one of the hosts to enter its bootloader.
.See *Remotely control blades via CMM* in the FortiAnalyzer-BigData Getting Started Guide for your appliance: [4500F](#) or [4500G](#).
For example: Go to *Blade System > Summary* and select Blade A2 to access the BMC (Blade Management Console).
3. Click the *BMC IPV4* link to enter the BMC for the host.
The default login credentials are on the Fortinet Product Credentials card
4. Go to *Remote Control > Console Redirection or iKVM/HTML5*.
5. Click *Power Control > Set Power Reset*.
6. Wait for the following options to appear. Use the arrow keys to select `bootloader` to open the bootloader's main page.



Bootloader Main Page

```
FortiAnalyzer-BigData-Bootloader 7.2.5
1). Configure Network.
2). Install OS.
3). Set Role. [Now Worker]
4). Set Chassis ID. [Now 1]
5). Set Blade ID. [Now 5]
6). Reset OS without Clearing User Data.
7). Reset OS and Clear User Data.
8). Upgrade Bootloader.
9). Extract hardware detailed information by lshw
10). Check and repair hard drives
11). Reset root password
0). Reboot.
Please input choice:
```

From the main page of the bootloader, you can select the following options:

- 1. Configure Network
- 2. Install OS
- 3. Set Role
- 4. Set Chassis ID
- 5. Set Blade ID
- 6. Reset OS
- 7. Reset OS and Clear User Data
- 8. Upgrade Bootloader
- 9. Extract hardware detailed information by lshw
- 10. Check and repair hard drives
- 11. Reset root password
- 12. Reboot
- sh. shell

1. Configure Network

The Configure Network option enables users to configure their IP, network mask, and network gateway information for the bootloader on the host in order to communicate with external servers hosting bootloader or FortiAnalyzer-BigData firmware images. Users can choose to specify static or DHCP IP addresses when available.



This option only configures the network for the bootloader, not the OS of the FortiAnalyzer-BigData host.

Before users can use this option to configure the network, they need to have the network interface associated with the external network. By default, the external network interface defaults to eth1.

```
Please input choice: 1
Please Choose Port:
eth0
eth1
Your Choice [eth1]:
Please Input IP/MASK [dhcp]: 10.106.2.168/24
Please Input Gateway [1]: 10.106.2.254
Your current input:
Device: [eth1]
IP/MASK: [10.106.2.168/24]
Gateway: [10.106.2.254]
Corrent? [Y/N/C]: Y_
```

2. Install OS

The Install OS option enables users to install FortiAnalyzer-BigData OS images on the host. Upon selection, users are prompted to provide server and image information. After confirmation, the FortiAnalyzer-BigData OS is downloaded from the server and installed.

Generally, users should use the `fazbdctl upgrade fazbd` command in FortiAnalyzer-BigData OS to upgrade the system software instead of using the bootloader Install OS option.

```
Please input choice: 2
Please choose method:
1). FTP
0). Cancel
Your choice: 1
Please input server IP [10.106.2.123]:
Please input file path [FAZBD.out]:
Please input username [ftp]:
Please input password:
Your current input:
Server IP: [10.106.2.123]
File path: [FAZBD.out]
Username: [ftp]
Password: []
Continue? [y/les/[n]o/[c]ancel: y_
```

3. Set Role

The Set Role option enables users to select a role for each host. You can see the current role of the host by the option.

In a FortiAnalyzer-BigData Security Event Manager architecture, each host has a designated role in order to collaborate with other hosts. There are two roles from the bootloader perspective: controller and worker.

- Controller: Refers to the Security Event Manager Controller and acts as the master of the other hosts.
- Worker: Nodes that are managed by the controller.

In a given cluster, only one active controller is allowed.

```
Please input choice: 3
1). Controller.
2). Worker.
Please choose blade role: 1_
```

4. Set Chassis ID

The Set Chassis ID is used to identify the chassis in multi-chassis cluster use case. Chassis IDs may range from 1 to 254. By default, it is 1. When you connect an extension chassis to an existing chassis cluster, the chassis ID needs to be changed to a unique number in 1 to 254 range. You can see the current Chassis ID by option.

```
Please input choice: 4
Please input chassis ID [1-254]: 1_
```

5. Set Blade ID

A Blade ID is used to identify the blade slot within a chassis. The order of the blade slots starts from the left side of the FortiAnalyzer-BigData appliance, starting from 1 to 14.

By default, all Blade IDs are set to reflect its physical slot number and users should not change the Blade ID. For example, the controller is in blade slot #2 and has a Blade ID of 2.

If you need to add a replacement blade to the chassis, you must first set the Blade ID to reflect its slot number so the firmware running on the blade knows its physical slot and its role.

```
Please input choice: 5
Please input blade ID [1-254]: 2_
```

6. Reset OS

The Reset OS option enables users to soft reset the FortiAnalyzer-BigData firmware of this BigData host. To soft reset the whole Security Event Manager, use `fazbdct1` CLI commands on the BigData Controller instead (see [Soft reset FortiAnalyzer-BigData on page 109](#)).



A soft reset only restores the firmware and will not touch the data volume.



If this action is performed on the BigData Controller, all the BigData member hosts will have to be rebooted during the progress in order to sync with the BigData Controller.

7. Reset OS and Clear User Data

The Reset OS and Clear User Data option enables users to hard reset the FortiAnalyzer-BigData firmware of this BigData host. To hard reset the whole Security Event Manager, use `fazbdctl` CLI commands on the BigData Controller instead (see [Hard reset FortiAnalyzer-BigData on page 110](#)).



This will restore the firmware AND clear all the data volume.

8. Upgrade Bootloader

The Upgrade Bootloader option enables users to specify server and image information to perform upgrades to the existing bootloader. When using this option, you must perform the step for all hosts individually.

Instead, it is recommended to upgrade the bootloader from the *Security Event Manager Controller* using the following command:

```
fazbdctl upgrade bootloader
```

This command allows you to upgrade the bootloader for all hosts at once.

9. Extract hardware detailed information by lshw

This command extracts hardware detailed information using the `lshw` command.

10. Check and repair hard drives

This command is used to verify and repair the integrity of file systems on storage devices.

11. Reset root password

This command resets the OS password of this host to the default "`fortinet@123`".

12. Reboot

The Reboot option enables you to reboot and restart the host.

sh. shell

If you enter `sh` into the Bootloader prompt, you can access the shell and use tools under `/sbin/`. For example, you can use `xfstool` to fix root disk errors if they occur.

General maintenance and best practices

To ensure that your FortiAnalyzer-BigData appliance runs smoothly, you need to perform regular maintenance tasks and follow best practices guidelines.

Backup and restore to external HDFS



For full instructions on how to backup and restore your data, see [Data backup on page 79](#) and [Data restore on page 84](#).

Schedule maintenance tasks for off-peak hours

Fortinet strongly recommends scheduling maintenance jobs for off-peak hours whenever possible, including jobs such as:

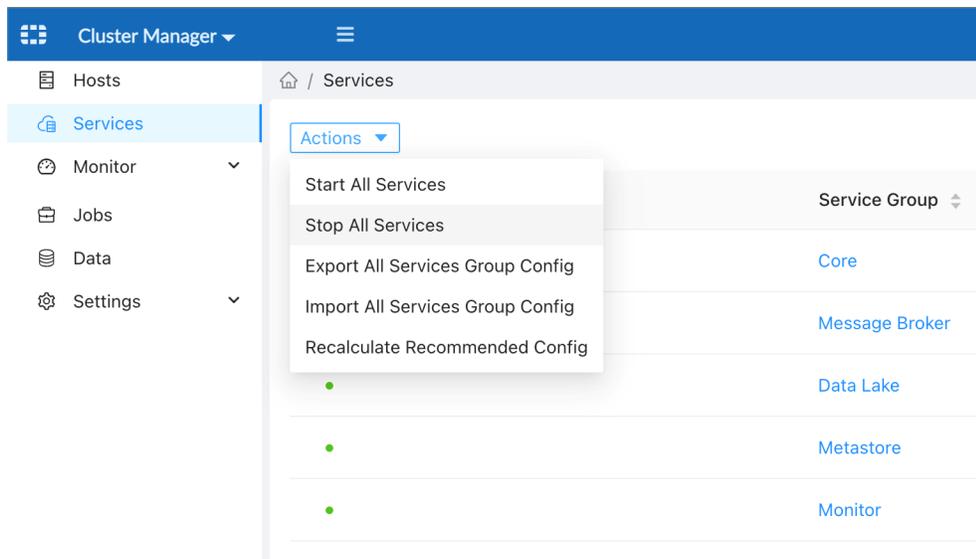
- Storage Pool Backup
- Data Rebalance

Graceful system shutdown

We recommend gracefully shutting down the FortiAnalyzer-BigData services before a planned power outage in order to help prevent data corruption in the stateful workload and conflicts in the distributed consensus.

To gracefully shut down the services:

1. Go to *Cluster Manager > Services*.
2. From the *Actions* dropdown, click *Stop All Services*.
Wait for the *Stop All Services* command to finish before powering off the chassis.



Maintain database integrity

To maintain database integrity, never power off a FortiAnalyzer-BigData unit without a graceful shutdown. Removing power without a proper shutdown can damage FortiAnalyzer-BigData databases.

Before removing power, always use the *Stop All Services* action from *Cluster Manager > Services > Actions*, or manually stop services in the following order:

1. Core
2. Message Broker
3. Data Lake
4. Metastore



After you power up your FortiAnalyzer-BigData unit again, you must manually select the *Start All Services* action from *Cluster Manager > Services > Actions* and make sure that all hosts, services and health checks are green before resuming system functions.



Fortinet strongly recommends connecting FortiAnalyzer-BigData units to an uninterruptible power supply (UPS) to prevent unexpected power issues that might damage internal databases.

Upgrade FortiAnalyzer-BigData

Before you upgrade FortiAnalyzer-BigData, ensure you have an FTP server that the FortiAnalyzer-BigData Security Event Manager Controller can access. Then put the FortiAnalyzer-BigData image on the FTP server.

Upgrade takes about 45 minutes. During the upgrade, the GUI is not available. Log collecting, LogView, and FortiView operations are also not available.



It is recommended to perform the upgrade via the GUI.

The upgrade process via the CLI may fail if the SSH connection is disrupted before the Controller hosts reboot during the early upgrade stage.



After the upgrade completes, proceed to upgrade bootloader as well. Your version of FortiAnalyzer-BigData bootloader must match your current version of FortiAnalyzer-BigData.

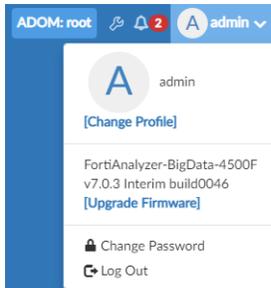
To upgrade the bootloader, see [Upgrade Bootloader](#).

See below for steps:

- [To upgrade FortiAnalyzer-BigData with the GUI:](#)
- [To upgrade FortiAnalyzer-BigData with the CLI: on page 101](#)

To upgrade FortiAnalyzer-BigData with the GUI:

1. Using the Management IP, sign into the FortiAnalyzer-BigData GUI.
2. In the banner, open the Account menu and click *Upgrade Firmware*.



3. Click *Upgrade* to access the System Upgrade dialog box.

System Upgrade

Select Source:

* Server IP:

User Name:

Password:

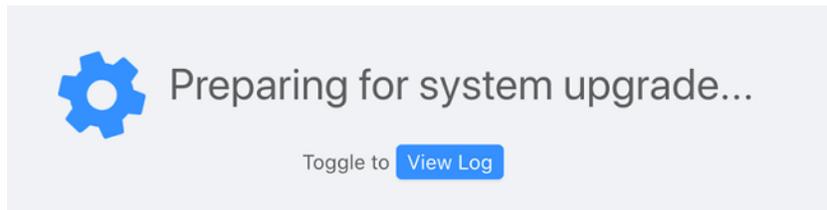
* File Path:

The system will be **rebooted** during the upgrade, and the whole process will take approximately 40 minutes. Users will not be able to access the system or use some of its features until the upgrade is fully complete. To view the progress, connect to the active Controller shell.

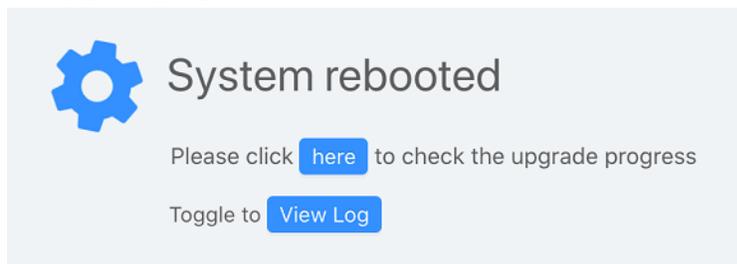
[Upgrade](#)

4. From the *Select Source* dropdown, select *FTP*, *SFTP*, *HTTP*, *HTTPS*, or *Upload File*.
5. Enter the source server's IP address, username, password, and file path. Leave the *User Name* and *Password* fields empty if they are not applicable.
6. Click *Upgrade*.

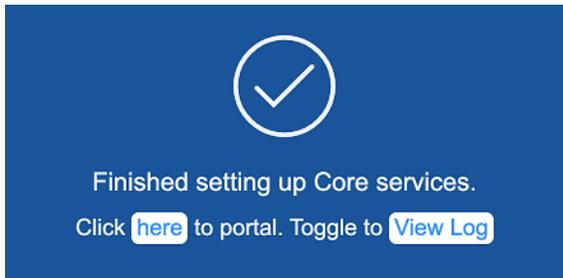
The system begins to prepare for the upgrade.



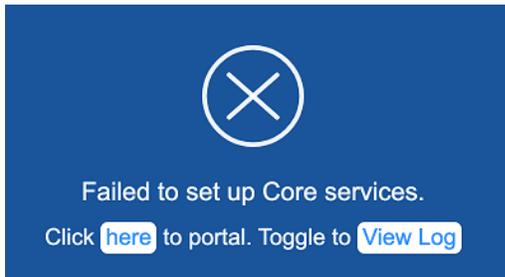
After the system finishes preparing, it loads a new page where you can see the current status and check the upgrade progress.



The upgrade takes about 45 minutes. If the upgrade is successful, you see the following message.



7. Click *here* to return to the FortiAnalyzer-BigData portal.
If the upgrade fails, you see the following message:



To troubleshoot the problem, see [What to do if an upgrade fails on page 118](#).

To upgrade FortiAnalyzer-BigData with the CLI:

You can also upgrade your FortiAnalyzer-BigData using the `fazbdctl` CLI command on the Security Event Manager Controller.

1. Access the Security Event Manager Controller CLI by establishing an SSH connection to the Cluster Management IP. See [Connect to the FortiAnalyzer-BigData CLI on page 12](#).



Starting the upgrade process from the Main Host instead of the Security Event Manager Controller may result in upgrade failure.

If this failure occurs, you must start again with a forced upgrade from the Security Event Manager Controller via SSH connection. Use the `-f` option to perform the forced upgrade. There is no forced upgrade option via the GUI.

2. Run the following command:

```
fazbdctl upgrade cluster -U <ftp_path> -u <user> -p <password>
```

Or, interactively,

```
fazbdctl upgrade cluster
```
3. Follow the onscreen instructions to select the source from FTP, SFTP, HTTP or HTTPS, and enter your server URL, upgrade file's zip file path, and FTP username and password.
The system upgrades the FortiAnalyzer-BigData Main Host and then the Security Event Manager. After a few minutes, the Security Event Manager Controller reboots.
4. After the Security Event Manager Controller reboots, reconnect to it and monitor the broadcast messages for progress.
5. Wait about 45 minutes for the following message to display on the terminal.
[100%] The system has been upgraded successfully.



Downgrading to prior versions of FortiAnalyzer-BigData is not supported.



When an upgrade is in progress, upgrades with the CLI will be blocked from execution with the following error message:

```
Error: Installation Process is already in progress
```

```
Lock ID: xxxxxxxxxx
```

```
To force unlock the installation process, use command: fazbdadm force-  
unlock-install xxxxxxxxxx
```

The force-unlock-install CLI command can be used to force release the upgrade process lock, but it should only be used when the upgrade process stops unexpectedly without releasing the process lock.

Scaling FortiAnalyzer-BigData

Scaling FortiAnalyzer-BigData:

You might need to scale the Security Event Manager of FortiAnalyzer-BigData by stacking multiple FortiAnalyzer-BigData appliances to add more storage and query throughput. For example, if you have an existing deployment and want more disk space to store logs for a longer period of time, you can scale out by adding one or more extender chassis. The log data as well as the computing and stateful workload will be distributed across all the hosts in the stacked appliances.

How to scale out

For FortiAnalyzer-BigData, you can scale out by adding more extender chassis. Stacking multiple FortiAnalyzer-BigData chassis enhances system redundancy and ensures high availability at the host level while expanding storage capacity and computational strength. For example, stacking two chassis enables an increased ingestion rate from 300000 logs/second to approximately 570000 logs/second.

Blade 1 (the Main hosts) on the two stacked chassis form HA in Active-Active mode. Both of these FortiAnalyzer-BigData blades will be able to receive logs. When stacked, FortiAnalyzer-BigData will also scale out Kudu tables to all tablets on all blades on two or more chassis.

Stacking 4500F units: The Main host OS operates on bare-metal blades so they cannot be managed by the Controller. Other hosts can be assigned using *Assign Role*, but Main host HA needs to be set up manually. After completing the steps to add an extender chassis, complete the steps in [Set up Main Host HA in stacked FortiAnalyzer-BigData 4500F chassis on page 104](#).

Stacking 4500G units: The Main host OS operates as containers, so they can be managed by the Controller the same as other services. Both the Main host and other host roles can be assigned using *Assign Role*, so you do not need to manually configure the Main host HA after adding an extender chassis.

To add an extender chassis:



The following operation removes all user data from the extender chassis.

-
1. Ensure both the main and extender chassis are running the same and the latest version of FortiAnalyzer-BigData. To do this, run the following command:

```
fazbdctl show version
```



Do not connect the links between both chassis until step 2 has been successfully completed; otherwise, it may cause an IP conflict and corruption in the distributed consensus.

2. On the extender chassis, access the *Security Event Manager Controller* and run the following command:

```
fazbdctl set appliance <chassis_id>
```

For information about accessing the Security Event Manager Controller, see [To connect to the Security Event Manager Controller: on page 12](#).



The chassis ID is an integer starting with 1, where 1 is the default chassis ID of the cluster formed by one chassis.

If you are adding the first extender to a cluster with one chassis, the *chassis_id* of the extender should be set to 2. You should increment this *chassis_id* as you stack more chassis to the cluster.

For example, the second extender's *chassis_id* should be set to 3 before you connect it to a cluster of two chassis. This function updates the chassis id on all hosts, and hard resets the FortiAnalyzer-BigData system.

Wait for the command to finish running without errors before you proceed with the next step.

3. **For 4500F units**, connect the 40GE links with QSFP between the Internal Switch Modules (Switch Module #1) of the extender and main chassis.
For 4500G units, connect the 100GE links with QSFP between the Internal Switch Modules (Switch Module #1) of the extender and main chassis.
4. On the main chassis, access the Security Event Manager Controller and run the following command to make sure the new hosts have been added:

```
fazbdctl show members
```


Wait until all the hosts' status shows as *Alive*. For 4500F units, there should be 13 additional hosts added as members. For 4500G units, there should be 14 additional hosts added as members.
5. Access the FortiAnalyzer-BigData GUI of the main chassis, and go to *Cluster Manager > Hosts*.
6. Click *Assign Role* to assign the newly added hosts.
New hosts should have a *New* label.
7. Wait for the *Assign* job to complete for all services to become healthy.
8. Follow the *Found New recommended configurations* notification on the *Cluster Manager > Hosts* view to apply the optimized configuration for the newly added hosts.
For 4500G units, the HA is automatically configured on the Main hosts after assigning the roles. For 4500F units, you must proceed to [Set up Main Host HA in stacked FortiAnalyzer-BigData 4500F chassis on page 104](#).

Set up Main Host HA in stacked FortiAnalyzer-BigData 4500F chassis

High Availability (HA) can be configured on the Main Hosts of stacked FortiAnalyzer-BigData 4500F chassis, ensuring system resilience by allowing for failover in the event of a Main Host failure. For more details, refer to [FortiAnalyzer Administration Guide](#).

For stacked 4500G chassis, the Main Host HA is automatically set up after assigning roles when scaling out. See [How to scale out on page 103](#).



Backup your configuration before proceeding to the next steps. For steps to backup configuration, see the [FortiAnalyzer Upgrade Guide](#).

To configure a Main Host HA cluster, go to *System Settings > HA* and configure the settings or change the cluster configuration. The configuration must be performed in the Main Hosts of the main and each extender chassis. The Main Host IP address must be set up for both the main and extender chassis to access the GUI and CLI.

There are two sections in *System Settings > HA*:

- *Cluster Status*: use this panel to monitor the Main Host HA cluster status.
- *Cluster Settings*: use this panel to create or change the Main Host HA cluster configuration. For more information, see [Cluster Settings options on page 105](#).

To set up Main Host HA in stacked FortiAnalyzer-BigData 4500F chassis:

1. Configure the Primary host for the Main Host HA cluster.
 - a. Connect to the GUI for the Main Host of the main chassis.
 - b. Go to *System Settings > HA*.
 - c. Configure the fields in *Cluster Settings* according to the descriptions in [Cluster Settings options on page 105](#).
 - d. Click *Apply* to save.
2. Configure the Secondary host(s) for the Main Host HA cluster.
 - a. Connect to the GUI for the Main Host of each extender chassis.
 - b. Go to *System Settings > HA*.
 - c. Configure the fields in *Cluster Settings* according to the descriptions in [Cluster Settings options on page 105](#).
 - d. Click *Apply* to save.
3. Monitor the *Cluster Status*. Wait for the *Configuration Sync* of the Secondary Role to be "In-Sync".

Cluster Settings options

Cluster Settings

Operation Mode	Select <i>Active-Active</i> to configure the Main Host on each chassis for HA. Select <i>Standalone</i> to stop operating in HA mode.
Preferred Role	Select the preferred role when this host first joins the Main Host HA cluster. If the preferred role is <i>Primary</i> , then this host becomes the primary host if it is configured first in a new Main Host HA cluster. If there is an existing primary host, then this host becomes a secondary host. The default is <i>Secondary</i> so that the host can synchronize with the primary host. A secondary host cannot become a primary host until it is synchronized with the current primary host.

Cluster Virtual IP

IP Address	The IP address for which the Main Host HA cluster is to provide redundancy.
Interface	The interface the Main Host HA cluster uses to provide redundancy.
Action	Click the plus (+) to add another virtual IP. Click the x to remove a virtual IP from the list.
Cluster Settings	
Peer IP	If this is the <i>Primary</i> , type the internal IP address (e.g., 198.18.2.1) of Main Host of the extender chassis . If this is a <i>Secondary</i> , type the internal IP address (e.g., 198.18.1.1) of Main Host of the main chassis .
Peer SN	Type the serial number of the Main Host for the chassis corresponding to the IP address used in the <i>Peer IP</i> .
Action	Click the plus (+) to add another Main Host in the cluster. Click the x to remove a Main Host from the cluster.
Group Name	Type a group name that uniquely identifies the Main Host HA cluster. All members in a cluster must have the same <i>Group Name</i> , <i>Group ID</i> , and <i>Password</i> .
Group ID	Type a group ID from 1 to 255 that uniquely identifies the Main Host HA cluster.
Password	A password for the Main Host HA cluster. All members of the Main Host HA cluster must have the same password.
Heart Beat Interval	The time the primary host waits between sending heartbeat packets, in seconds. The heartbeat interval is also the amount of time that secondary hosts wait before expecting to receive a heartbeat packet from the primary host. By default, the <i>Heart Beat Interval</i> is set to 4.
Heart Beat Interface	Select the interface used to send heartbeat packets.
Failover Threshold	The number of seconds that one of the cluster hosts waits to receive HA heartbeat packets from other cluster hosts before assuming that the other cluster hosts have failed. This value corresponds to <i>Heart Beat Interval</i> x 3 and it is automatically updated based on the configured <i>Heart Beat Interval</i> . For example, the failure is detected after 12 seconds with the default settings: <ul style="list-style-type: none"> • <i>Heart Beat Interval</i>: 4 • <i>Failover Threshold</i>: 12 The <i>Heart Beat Interval</i> can be increased or decreased to adapt to latency conditions of your network and to detect legitimate failures.
Priority	The priority or seniority of the secondary host in the cluster.
Log Data Sync	This option is on by default. It provides real-time log synchronization among the peers in the cluster.

How to remove a chassis from a stacked setup

If you have an established multi-chassis FortiAnalyzer-BigData Security Event Manager and need to scale down, you can remove any extender chassis you have added to the main chassis.



Removing an extender chassis will hard reset BOTH the extender and the main chassis. All user data and configurations will be lost. The entire process takes approximately an hour.

Remove an extender chassis

To remove an extender chassis:



Perform this process in a location where you can immediately disconnect the Internal Switch Modules at step 3. This step must be done in a timely manner.

1. Stop the DHCP service. To do this, access the *Security Event Manager Controller* of the main chassis and run the following command:

```
systemctl stop dhcpd
```

For information about accessing the Security Event Manager Controller, see [To connect to the Security Event Manager Controller: on page 12](#).
2. Reset all hosts connected to the current controller by running the following command:

```
fazbdctl unstack-chassis
```

Wait for the message "[Info] factory reset members done, and [info] factory reset the controller" in the command output before proceeding to the next step without any delay.
3. Disconnect the connection between the *Internal Switch Modules* (Switch Module #1) between the main and extender chassis.
4. On the main chassis, access the *Security Event Manager Controller*.
 - a. Access the Bootloader of the active controller. See [Bootloader on page 91](#).
 - b. From Bootloader, select option 7, *Reset OS and Clear User Data*, to hard reset the host.
 - c. Wait until the Bootloader finishes rebooting and log back into the controller.
 - d. Initialize the cluster using the following command:

```
fazbdctl init cluster
```
 - e. Configure the management IP and main host IP.
For more information about setting up the network, see the Getting Started Guide for your device on the [Fortinet Document Library](#).
5. On the extender chassis, access Blade A2.
 - a. Access the Bootloader of the Blade A2. See [Bootloader on page 91](#).
 - b. From Bootloader, select option 3, *Set Role, input choice: Controller*, to set the role to *Controller*.
 - c. From Bootloader, select option 7, *Reset OS and Clear User Data*, to hard reset the host.
 - d. Wait until the Bootloader finishes rebooting and log back into the controller.

- e. Initialize the cluster by running the following command:

```
fazbdctl init cluster
```

- f. Configure the management IP and main host IP.

For more information about setting up the network, see the Getting Started Guide for your device on the [Fortinet Document Library](#).

Reset FortiAnalyzer-BigData

This section contains information on how to reset FortiAnalyzer-BigData. There are two ways to perform a reset:

- [Soft reset FortiAnalyzer-BigData on page 109](#)
- [Hard reset FortiAnalyzer-BigData on page 110](#)

Soft reset FortiAnalyzer-BigData

You can try to soft reset your FortiAnalyzer-BigData Security Event Manager to recover from a system level failure. This process takes about 45 minutes.

Soft reset does the following:

- Reset the OS partition on each of the hosts while keeping all data volume intact.
- Reset the FortiAnalyzer-BigData power.
- Align all the host OS.

To soft reset FortiAnalyzer-BigData:



Before proceeding with the steps below, your version of FortiAnalyzer-BigData bootloader must match your current version of FortiAnalyzer-BigData. Check the version of your bootloader and upgrade it to match your FortiAnalyzer-BigData as needed.

To check the bootloader version, run the following command from the *Security Event Manager Controller*: `fazbdctl show version`

To upgrade the bootloader, see [Upgrade Bootloader](#).

1. Access the Security Event Manager Controller (see [To connect to the Security Event Manager Controller: on page 12](#)) and run the following command:
`fazbdctl reset cluster`
For more information and additional CLI options, see the reset command in the CLI Reference in the [Fortinet Doc Library](#).
The Security Event Manager Controller will reboot after a few minutes.
2. Reconnect to the Security Event Manager Controller after it reboots and monitor the broadcast messages for progress.
3. Wait about 45 minutes until the following message is displayed on the terminal:
[100%] The system has been upgraded successfully.

Hard reset FortiAnalyzer-BigData



Improperly resetting your FortiAnalyzer-BigData may result in losing all data.

When you hard reset your device, the command resets the OS on each host and formats all data drives. All log data and configurations will be lost. FortiAnalyzer-BigData shuts down during the reset process. The entire process takes approximately 45 minutes.

You can add an extra option to the reset command to keep certain configurations constant:

- `all-settings` resets all settings.
- `all-except-ip` keeps the public IP constant
- `all-except-ssh` keeps the ssh public key constant.
- `all-except-ip-ssh` keeps the ssh public key and public IP constant.

For more information about extra CLI options, see the reset command in the CLI Reference in the [Fortinet Doc Library](#).

To hard reset your FortiAnalyzer-BigData:



Before proceeding with the steps below, your version of FortiAnalyzer-BigData bootloader must match your current version of FortiAnalyzer-BigData. Check the version of your bootloader and upgrade it to match your FortiAnalyzer-BigData as needed.

To check the bootloader version, run the following command from the *Security Event Manager Controller*: `fazbdctl show version`

To upgrade the bootloader, see [Upgrade Bootloader](#).

1. Access the Security Event Manager Controller, and run the following command:
`fazbdctl reset cluster [--all-settings|--all-except-ip|--all-except-ssh|--all-except-ip-ssh]`
The Security Event Manager Controller reboots after a few minutes.
2. After the Security Event Manager Controller reboots, re-connect to it and run the following command to verify that all members are detected and that the version is up-to-date:
`fazbdctl show members`
3. After verifying that all the members have a *Joined* state and status is not failed, run the following command to initialize the Security Event Manager:
`fazbdctl init cluster`
4. Wait about 45 minutes until the following message is displayed on the terminal:
[100%] The system has been initialized successfully.

Data-at-rest encryption on FortiAnalyzer-BigData 4500F

The FortiAnalyzer-BigData 4500F offers data-at-rest encryption support to enhance the security of log data stored in the system. By leveraging Linux Unified Key Setup (LUKS) and dm-crypt, it encrypts all data disk partitions on the cluster hosts while no changes are required on the application logic or schema. This ensures that the log data remains protected and inaccessible even if unauthorized access to the physical storage is obtained. This encryption is managed through a cluster-level passphrase that provides consistency and simplifies administration.

Consider the following limitations to make an informed decision regarding the implementation of data-at-rest encryption in your FortiAnalyzer-BigData 4500F:

- The Main host blade does not undergo encryption, which means that raw logs may be buffered in the FAZ blade for a short period.
- Only the data disks of the Security Event Manager cluster hosts can be encrypted. The OS boot drive, which solely contains binaries, scripts, and server logs, is not encrypted.
- Enabling encryption necessitates re-formatting of the partitions. This process can only be performed during a fresh installation or a hard reset.
- No reversal of encryption once enabled. Once encryption is enabled, there is no way to revert it unless a factory reset is performed.



We strongly suggest you save the passphrase in a secure place. If the passphrase is lost and the encrypted volume is inactive (luks closed), data in the encrypted volumes will no longer be accessible.



Enabling data-at-rest encryption may lead to a performance degradation: a 17% slower disk I/O performance based on our lab test. However, the overall impact on performance may not be significant (< 10%) as most workloads are not heavily dependent on disk I/O.

Initialize data-at-rest encryption

Data-at-rest encryption can only be performed during a fresh installation or after a factory reset.

To initialize data-at-rest encryption as part of a factory reset, you must execute the `init` command with `--encrypt-data-disks`. See the steps below.

To initialize data-at-rest encryption as part of a factory reset:

1. Access the Security Event Manager Controller, and run the following command:
`fazbdctl reset cluster [--all-settings|--all-except-ip|--all-except-ssh|--all-except-ip-ssh]`
The Security Event Manager Controller reboots after a few minutes.
2. After the Security Event Manager Controller reboots, re-connect to it and run the following command to verify that all members are detected and that the version is up-to-date:
`fazbdctl show members`
3. After verifying that all the members have a *Joined* state and status is not failed, run the following command:
`fazbdctl init cluster --encrypt-data-disks`
4. Follow the prompt to set a passphrase. Record the passphrase in a secure place and follow the remaining reset steps to complete initializing the cluster.
See [Hard reset FortiAnalyzer-BigData on page 110](#).

Open encrypted data disk partitions

This topic contains examples of opening the encrypted data disk partitions in the following scenarios:

- [After a blade power cycle](#)
- [After a graceful chassis power cycle](#)
- [When performing an upgrade or soft reset](#)
- [When scaling out the cluster or replacing a blade](#)

After a blade power cycle

Every time the FortiAnalyzer-BigData OS boots up after a planned or unplanned power cycle, you need to open the encrypted data disk partitions with your passphrase.

There are two methods to open them.

To open the encrypted data disk partitions via the BMC remote console:

1. Access the BMC remote console of the restarted blade.
2. Log in to the OS using root credentials.
3. Follow the prompt to enter the cluster encryption passphrase for opening the data disk partitions.

```
CentOS Linux 7 (Core)
Kernel 5.4.200-1.el7.elrepo.x86_64 on an x86_64

blade-198-18-1-13 login: root
Password:
Last login: Wed Jul 26 00:05:40 on tty1
Enter Encrypted Volumes Passphrase to Mount (Tab to flip character mask):*****
/data1 mounted
/metadata1 mounted
/data2 mounted
/metadata2 mounted
[root@blade-198-18-1-13 ~]# _
```

To open the encrypted data disk partitions via the Controller shell:

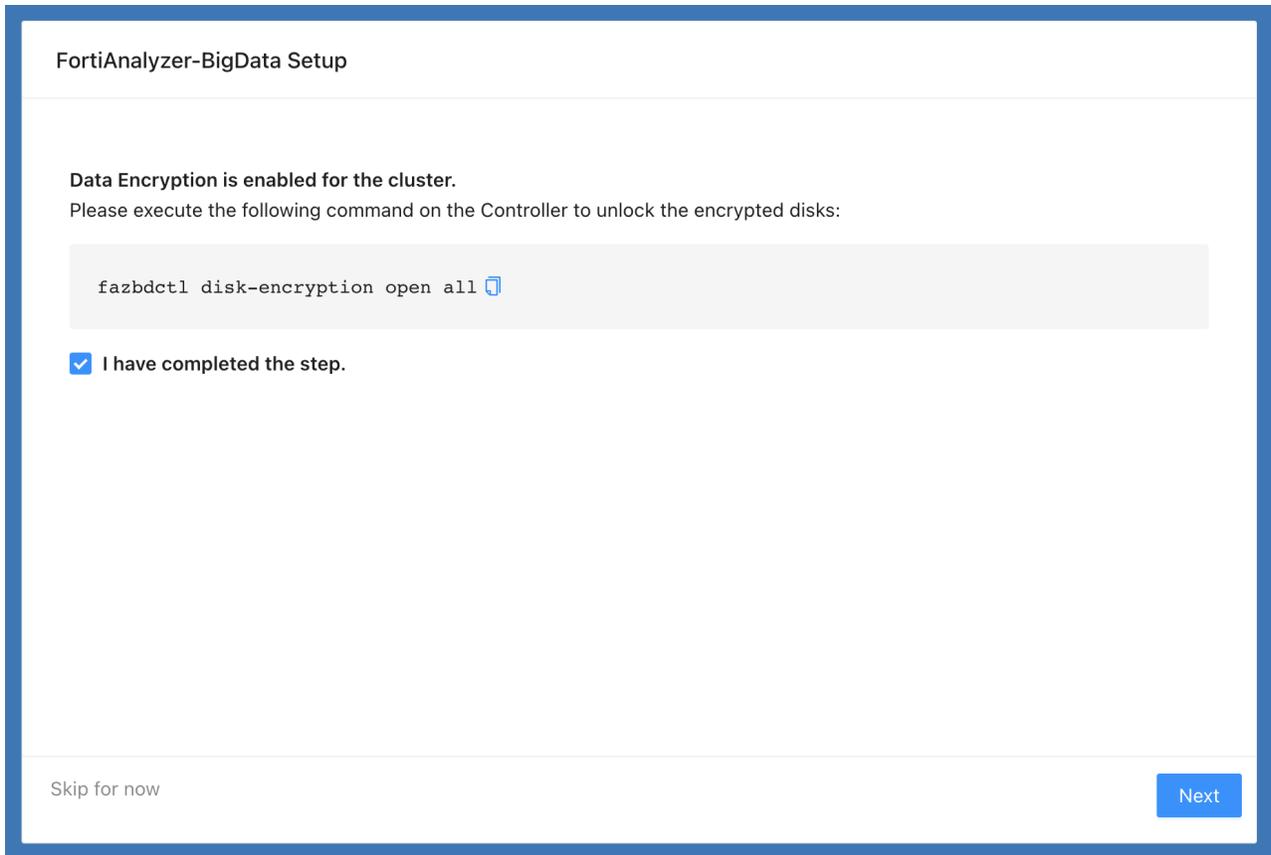
1. Connect to the Controller shell.
2. Utilize the `fazbdctl disk-encryption open {host}` command to open the encrypted disk partitions on the restarted blade.

```
[root@blade-198-18-1-2 ~]# fazbdctl disk-encryption open blade-198-18-1-5
The process will take several minutes. Do not power off the blades during the progress. Are you sure? [Y]es/[N]o/[C]ancel : y
Enter the current passphrase for cluster data encryption: *****
validating the passphrase...
validated
opening data disks on blade-198-18-1-5...
INFO: LUKS open nvme0n1p1 on /dev/nvme0n1p1
INFO: LUKS open nvme0n1p2 on /dev/nvme0n1p2
INFO: LUKS open nvme1n1p1 on /dev/nvme1n1p1
INFO: LUKS open nvme1n1p2 on /dev/nvme1n1p2
INFO: Mount crypto luks device /dev/mapper/nvme0n1p1
INFO: Mount crypto luks device /dev/mapper/nvme0n1p2
INFO: Mount crypto luks device /dev/mapper/nvme1n1p1
INFO: Mount crypto luks device /dev/mapper/nvme1n1p2
Done: opened the encrypted data disks on the requested host(s).
```

After a graceful chassis power cycle

To open the encrypted data disk partitions after a graceful chassis power cycle:

1. Follow steps to gracefully shut down services in the GUI.
Go to *Cluster Manager > Services > Actions > Stop All Services*, and then power off the blades.
See steps to perform a graceful system shutdown in [General maintenance and best practices on page 97](#).
2. Once the chassis (including all blades) is powered on again, the web GUI will display the Setup Wizard upon opening.



3. Follow the instructions to connect to the Controller shell and open the encrypted disk partitions. See [To open the encrypted data disk partitions via the Controller shell: on page 113.](#)
4. Once the data disk partitions are opened, proceed with the Setup Wizard GUI to start all services.

When performing an upgrade or soft reset

To open the encrypted data disk partitions during an upgrade or soft reset:

1. Follow steps to perform a system upgrade or soft reset.
For steps to upgrade, see [Upgrade FortiAnalyzer-BigData on page 99.](#)
For steps to soft reset, see [Soft reset FortiAnalyzer-BigData on page 109.](#)
2. After completing the pre-upgrade phases, the Controller will reboot and start to sync all members.
Once the members have re-joined, the upgrade process will pause for you to open the encrypted data disk partitions.
3. To proceed, reconnect to the Controller shell and follow the message prompts to open the encrypted disk partitions on all cluster hosts.
Once the encrypted disk partitions are opened, the upgrade will continue automatically.

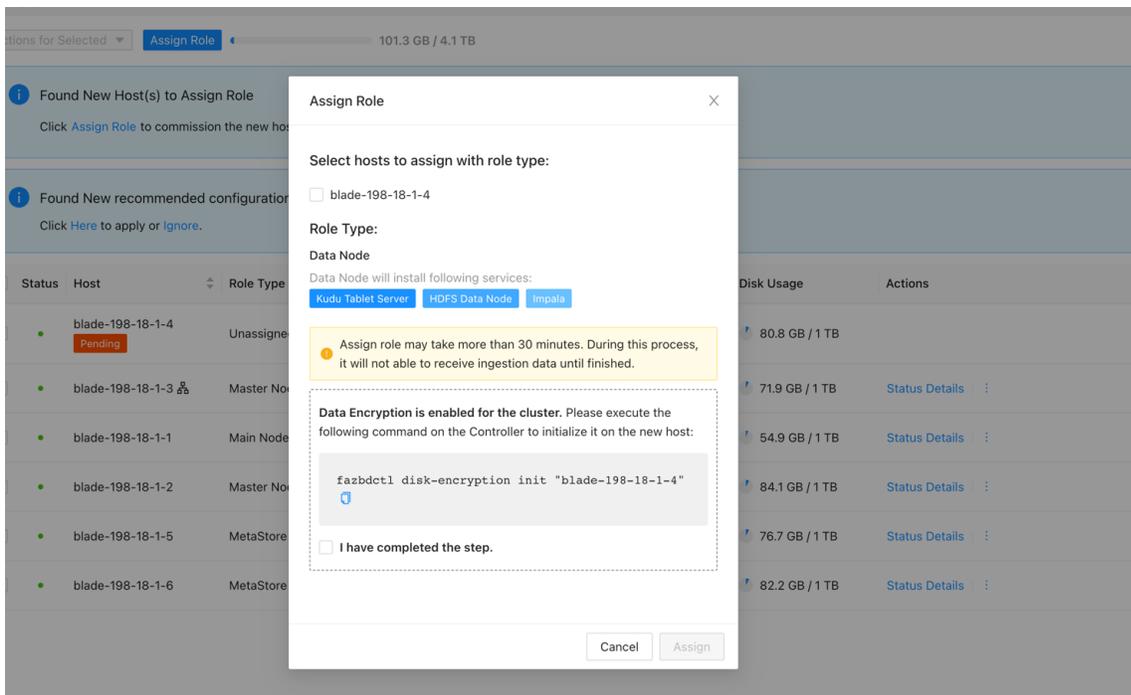
```
[40%] Upgrade/Reset progress has been paused due to data encryption. Please execute [fazbdctl disk-encryption open all] to continue the progress.

[root@blade-198-18-1-2 ~]$ fazbdctl disk-encryption open all
The process will take several minutes. Do not power off the blades during the progress. Are you sure? [Y]es/[N]o/[C]ancel : y
Enter the current passphrase for cluster data encryption: *****
validating the passphrase...
validated
opening encrypted data disks on 198.18.1.10...
opening encrypted data disks on 198.18.1.11...
opening encrypted data disks on 198.18.1.12...
opening encrypted data disks on 198.18.1.13...
opening encrypted data disks on 198.18.1.2...
opening encrypted data disks on 198.18.1.3...
opening encrypted data disks on 198.18.1.4...
opening encrypted data disks on 198.18.1.5...
opening encrypted data disks on 198.18.1.6...
opening encrypted data disks on 198.18.1.7...
opening encrypted data disks on 198.18.1.8...
opening encrypted data disks on 198.18.1.9...
```

When scaling out the cluster or replacing a blade

To open the encrypted data disk partitions when scaling out or replacing a blade:

1. Follow steps to scale out or replace a blade.
 For steps to scale out, see [How to scale out on page 103](#).
 For steps to replace a blade, see [How to replace a blade on a FortiAnalyzer-BigData appliance on page 120](#).
2. During the step to Assign Role for the blade, a confirmation prompt will appear in the GUI. This prompt will provide instructions on how to initialize data encryption on the new or replaced host(s).



3. Follow the instructions provided in the prompt to connect to the Controller shell and initialize the encrypted disk partitions on the new or replacement host(s). For example, connect to the Controller shell and enter the following command: `fazbdctl disk-encryption init {host}`.
 Once completed, you can proceed with the Assign Role process.



If the step to initialize the encrypted disk partitions is missed, the Assign Role task will fail during the pre-flight check.

Troubleshooting

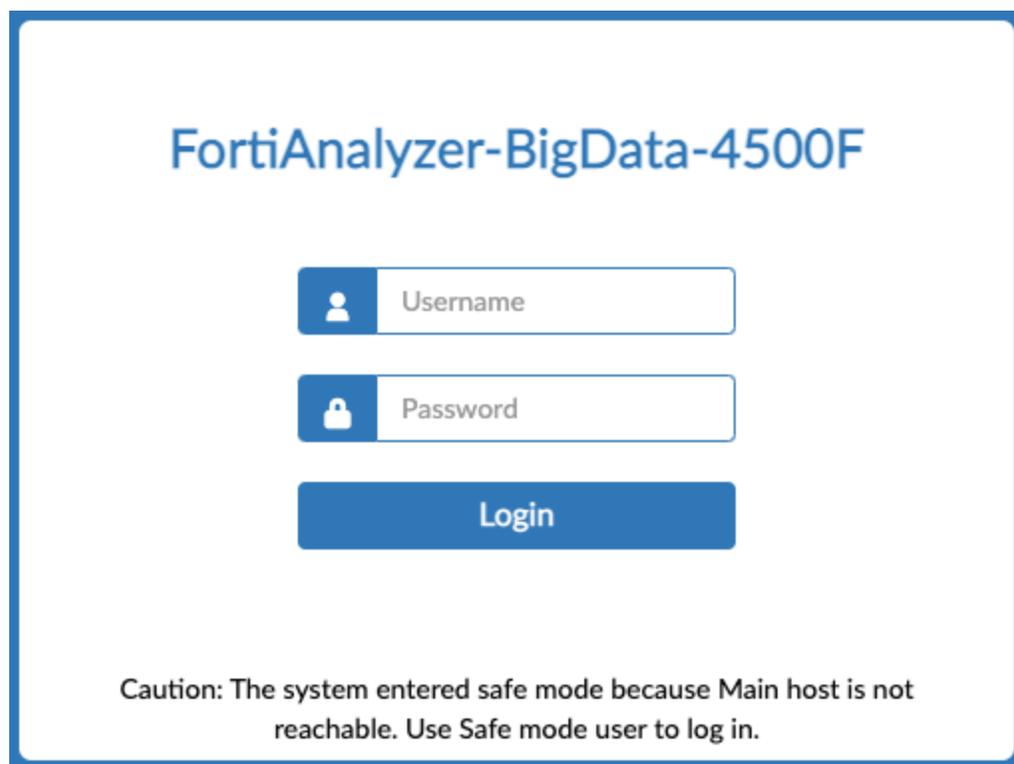
This section contains troubleshooting tips for issues you might encounter when working with FortiAnalyzer-BigData.

Safe Mode

The FortiAnalyzer-BigData web GUI will enter *Safe Mode* when the Main host becomes unhealthy or unreachable from the Security Event Manager. This can be caused by a loss of power to the Main host, hardware failure, network segregation, Main host OS crash, etc. Once in Safe Mode, only [Global Search](#) and [Cluster Manager](#) are accessible. The other tiles on the home page are grayed out.

To log in to FortiAnalyzer-BigData in Safe Mode, enter the Username admin and no password. You will be prompted to change the password the first time you log in.

For example, see the image below from a FortiAnalyzer-BigData 4500F GUI.



FortiAnalyzer-BigData will exit Safe Mode once the Main host is healthy again.

What to do if an upgrade fails

An upgrade might fail with the following error conditions:

Error condition	Troubleshooting suggestion
An error message displaying: <ul style="list-style-type: none"> "get image failed" "could not find image" 	Make sure image from the hosting server is accessible.
An error message displaying: <ul style="list-style-type: none"> "checksum verification failed" 	Check the image file integrity.
The Security Event Manager Controller cannot boot up after an upgrade and you cannot connect to the Security Event Manager Controller	Perform the following steps: <ol style="list-style-type: none"> 1. Access the bootloader of the Security Event Manager Controller (see Bootloader on page 91). 2. Select the "Backup" option to restore the last working OS image to the system.

You can also retry a failed upgrade by using the option flag -o in the upgrade command.

Enter the following command to retry upgrading from where it fails:

```
fazbdctl upgrade cluster -o retry
```

What to do if a soft reset fails

A soft reset might fail with the following error conditions:

Error condition	Troubleshooting suggestion
An error message displaying: <ul style="list-style-type: none"> "checksum verification failed" "could not find image" 	Perform an upgrade with the image of the intended version or latest version.
The Security Event Manager Controller cannot boot up after an upgrade and you cannot connect to the Security Event Manager Controller	Perform the following steps: <ol style="list-style-type: none"> 1. Access the bootloader of the Security Event Manager Controller (see Bootloader on page 91). 2. Select the "Backup" option to restore the last working OS image to the system. 3. Access the Security Event Manager Controller and perform an upgrade via fazbdctl CLI commands (see Upgrade FortiAnalyzer-BigData on page 99) with the image of the intended version or latest version. 4. Rerun the reset command to perform a soft reset.

You can also retry a soft reset by using the option flag (-o) in the rest command:

- Enter the following command to retry soft reset on the cluster from where it fails:
`fazbdctl reset cluster -o retry`
- Or enter the following command to retry soft reset on the cluster from the very beginning:
`fazbdctl reset cluster -o restart`

What to do if a hard reset fails

A hard reset might fail with the following error conditions:

Error condition	Troubleshooting suggestion
The reset failed to complete before the Security Event Manager Controller reboots	Upgrade the system to latest version (see Upgrade FortiAnalyzer-BigData on page 99) and then try resetting again.
The Security Event Manager Controller cannot start or the system is not accessible after a hard reset.	Perform the following steps: <ol style="list-style-type: none"> 1. Access the bootloader of the Security Event Manager Controller (see Bootloader on page 91). 2. Select the "Backup" option to restore the last working OS image to the system. 3. Access the Security Event Manager Controller and perform an upgrade via fazbdctl CLI commands with the image of the intended version or latest version. 4. Rerun the reset command to perform a hard reset.

How to repair disk failures

If you see a "disk failure" message in any system logs, it might indicate that the FortiAnalyzer-BigData is experiencing hard disk issues. You can try to repair these issues using software methods.

To repair disk failure issues:

1. Access the bootloader of the host that has disk failure symptoms (see [Bootloader on page 91](#)).
2. From the bootloader, enter `sh` to enter the shell.
3. In the shell, run `xfs_repair` to fix the hard disk issue.

If the problem persists after running the software fix, you may need to replace the hard disk.

How to replace a blade on a FortiAnalyzer-BigData appliance

This section contains instructions on how to gracefully remove and replace a malfunctioning hardware blade running one of the hosts in an active system. In order to allow the high availability mechanism to take effect, only one blade can be decommissioned at a time.

Finding a blade's location



A blade's host name uses the following naming convention:

`blade-198-18- $\{chassis\ ID\}$ - $\{blade\ ID\}$`

For example, in a blade named *blade-198-18-1-3*, the number *1* represents the chassis ID and *3* represents the blade ID. Therefore, the blade is the third blade to the left on the first chassis. The internal IP of the blade is *198.18.1.3*.

To replace a blade:



Before proceeding with the steps below, your version of FortiAnalyzer-BigData bootloader must match your current version of FortiAnalyzer-BigData. Check the version of your bootloader and upgrade it to match your FortiAnalyzer-BigData as needed.

To check the bootloader version, run the following command from the *Security Event Manager Controller*: `fazbdctl show version`

To upgrade the bootloader, see [Upgrade Bootloader](#).

1. Power off the malfunctioned blade, and then remove the blade from the chassis.
2. From *Security Event Manager Controller* run the following command to enable installation of system software on the new blade: `fazbdctl decommission <host name or blade IP of the blade to replace>`
3. Insert the replacement blade, and power it on.
4. Monitor the status of the blade from CMM. See *Remotely control blades via CMM* in the FortiAnalyzer-BigData Getting Started Guide for your appliance: [4500F](#) or [4500G](#).
Wait until the host OS completes booting. This may take 5-10 minutes.
5. After the host boots up and joins the cluster, it will appear in *Cluster Manager > Hosts* web GUI.
6. Go to *Cluster Manager > Hosts* and click *Assign Role* to recover the host. The new host will have a *pending* status.
7. Wait for the *Assign job* to complete for all services to become healthy.
8. Follow the *Found New* recommended configurations notification in the *Cluster Manager > Hosts* view (if it appears) to apply the optimized configuration for the newly added hosts.

How to reset a single host

This section contains instructions on how to gracefully reset a software malfunctioned host in a running system. In order to allow the high availability mechanism to take effect, only one host can be reset at a time.



Finding a blade's location

A blade's host name uses the following naming convention:

blade-198-18-`{chassis ID}`-`{blade ID}`

For example, in a blade named *blade-198-18-1-3*, the number *1* represents the chassis ID and *3* represents the blade ID. Therefore, the blade is the third blade to the left on the first chassis. The internal IP of the blade is *198.18.1.3*.

To reset a single host:

1. Access the bootloader of the malfunctioned host (see [Bootloader on page 91](#)), enter the Reset OS option and wait until it finishes rebooting.
2. Monitor the status of the blade or host.
 - a. Monitor the status of the blade from the DMM. See *Remotely control blades via CMM* in the FortiAnalyzer-BigData Getting Started Guide for your appliance: [4500F](#) or [4500G](#).
Wait until the host OS completes booting. This may take 5-10 minutes.



Cluster failover

When resetting the *Controller* host, allow up to 15 minutes for the failover mechanism to take effect. Once the mechanism is in effect, the *Security Event Manager IP* and *Cluster Manager* can be accessed with the GUI.

3. After the host reboots and joins the cluster, it will appear in *Cluster Manager > Hosts*.
4. From the *Hosts* page, click *Assign Role* to recover the role on the host. The new host should display a *pending* label. When the *Assign Role* job is complete, the host reset is done.

How to rebalance the data

This topic contains instructions on how to rebalance the data across the *Security Event Manager* hosts. The data is balanced automatically by default. However, in the circumstances such as a host failure, reset, or replacement, data may get skewed. The *Data balance Check* in *Cluster Manager > Monitor > Health* periodically checks for data balance and fails if the data skews high. The built-in *Data Rebalance (All)* job can be used to rebalance the data.



Maintenance window

By default, the data will be unavailable during the run to ensure data integrity. Choose a maintenance window when the log receiving rate is low.

To rebalance the data:

1. Go to *Cluster Manager > Jobs* and click *Create Custom Job*.
2. Select *Data Rebalance (All)* and *Schedule "Manual"*, then click *Create*.
3. In the *Jobs* table, locate the *Data Rebalance (All)* row, and click *Run in Actions*. Allow approximately 1 hour for the job to execute, depending on the data size.
4. After the job is finished, go to *Cluster Manager > Monitor > Health*. If the *Data balance Check* status is *Failed*, click *Run Test* to rerun the test. Ensure the test status is *Success*.

How to recover from an unhealthy service status

The services in the Security Event Manager are highly available and fault tolerant with data are replicated three times into different data hosts. If any one of the BigData hosts goes down, you can expect some service degradation (such as dropped insert rate and query performance), but all basic functionalities (such as FortiView, and LogView) are preserved with no data loss. While the system is mostly self-healing from failures, manual operation is required to address certain failure incidents.

The Monitor page contains tools to help you monitor the status and health of the hosts and services (see [Monitor on page 26](#)). We suggest scheduling a routine monitoring and maintenance window, and set up system alerts to enable rapid remediations and fault prevention. If you need to shut down your FortiAnalyzer-BigData, follow the best practices (see [General maintenance and best practices on page 97](#) to avoid damaging your database.

Stateful workloads occasionally require manual responses to recover from incidents. When unhealthy workloads are detected, check the status of all BigData hosts to ensure they are all functioning. In general, you should address host-level incidents first before going into the service level.

The following sections contain troubleshooting tips for when FortiAnalyzer-BigData services have an unhealthy status.

Core services

Core / Query

If Query service is unhealthy, or if FortiView or LogView stops working, you can try the following:

1. From *Cluster Manager > Services*, check if the Data Lake service group is healthy, if not, fix it first.
2. From *Cluster Manager > Services > Core*, manually restart the Query service, and then wait a few minutes to see if the issue is fixed.

Core / Ingestion

If the Ingestion service is unhealthy, or if the log insert rate remains at zero while receiving rate is higher, you can try the following:

1. From *Cluster Manager > Services*, check if the Message Broker service group is healthy, if not, fix it first.
2. In *Cluster Manager > Services > Core*, manually *Restart* the Ingestion service, and then wait a few minutes to see if the issue is fixed.
3. If the issue persists after the restart, go to *Cluster Manager > Jobs > Create Custom Job*, and select *Kafka Deep Clean* as the template.
4. Find the newly created "Kafka Deep Clean" job in the job list and click *Run*.



This will purge all the data in the queue and start a fresh Pipeline. Any unprocessed data will be lost.

Core / Pipeline

If the Pipeline service is unhealthy, or if the Pipeline Health Check in *Monitor > Health* remains unhealthy for hours, you can try the following:

1. In *Cluster Manager > Services*, check if the Data Lake and Message Broker service groups are healthy, if not, fix them first.
2. In *Cluster Manager > Services > Core*, manually restart the Pipeline service, and then wait a few minutes to see if the issue is fixed.
3. If the issue persists after a few hours, go to *Cluster Manager > Jobs > Create Custom Job* and select *Purge Data Pipeline* as the template.
4. Find the newly created "Purge Data Pipeline" job in the job list and click *Run*.



This will purge all the data in the queue and start a fresh Pipeline. Any processed data will be lost.

Data Lake services

Data Lake / Impala

If the Impala service is unhealthy, you can try the following:

1. Check if the Metastore service group is healthy, if not, fix it first.
2. From *Cluster Manager > Services > Data Lake*, manually *Restart* the Impala service and wait a few minutes to see if the issue is fixed.

Data Lake / Kudu

If the Kudu service is unhealthy, you can try the following:

1. From *Cluster Manager > Services*, manually *Stop* the Core service group.
2. Check if the Metastore service group is healthy, if not, fix it first.
3. From *Cluster Manager > Services > Data Lake*, manually *Restart* the Kudu service and wait a few minutes to see if the issue is fixed.

4. If the issue persists after the restart and the log indicates that Kudu failed to synchronize time, go to *Cluster Manager > Jobs > Create Custom Job* and select *NTP Sync* as the template.
5. Find the newly created *NTP Sync* job in the job list and click *Run*.
6. After the job finishes running, manually *Start* the Kudu service again to see if the status becomes healthy.
7. Once the Kudu service is healthy, manually *Start* the Core service group again.

If the Database Health Check in *Monitor > Health* remains unhealthy for hours but the Kudu service status is healthy, you can try the following:



The Database Health Check may temporarily fail when the Storage Pool Restore or Data Rebalance job is running. Once the jobs are finished running, the status will automatically clear. Make sure those jobs are not running before troubleshooting.

1. From *Cluster Manager > Services*, manually *Stop* the Core service group.
2. Wait about 15 minutes and then navigate to *Monitor > Health* to rerun the Database Health Check.
3. If the health check returns as healthy, return to the Services page to manually *Start* the Core service group.

Message Broker services

Message Broker / Kafka

1. If the Kafka service is unhealthy, you can try the following:
2. From *Cluster Manager > Services*, manually *Stop* the Core service group.
3. Go to *Cluster Manager > Services > Message Broker*, and manually *Restart* the Kafka service and check that the status becomes healthy.
4. If the issue remains after the restart, go to *Cluster Manager > Jobs > Create Custom Job* and select *Kafka Deep Clean* as template.
5. Find the newly created "Kafka Deep Clean" job in the job list and click *Run*.



This will purge all the data in the queue and start a fresh Pipeline. Any processed data will be lost.

6. Return to *Cluster Manager > Services* and manually *Start* the Core service group.

Metastore / HDFS

If the HDFS service is unhealthy, or if the HDFS related Health Checks in *Monitor > Health* are remains, you can try the following:

1. From *Cluster Manager > Services > Metastore*, manually *Restart* the HDFS service, and then wait a few minutes to see if the status changes to healthy.
2. If the issue persists after restart and the logs indicate the HDFS is in safe mode, go to *Cluster Manager > Jobs > Create Custom Job* and select *HDFS Safemode Leave* as the template.
3. Find the newly created "HDFS Safemode Leave" job in the job list and click *Run*.

How to recover from a full disk

The FortiAnalyzer-BigData data life cycle can be managed via Cluster Manager GUI (see [Manage storage policy on page 78](#)). If the data disk on your hosts begin to reach full capacity and are causing the Data Lake services to become unhealthy, you can do the following:

1. From *Cluster Manager > Services*, manually *Stop* the Core service group except the catalog service. The catalog service needs to continue running.
2. Go to *Cluster Manager > Data*, expand the Root Storage Pool and click *Action > Manage Data Lifecycle*.
3. In the *Maximum Age* field, reduce the number of days for storing data and click *OK*.
4. Go to *Cluster Manager > Jobs*, and locate and *Run* the Data Retention job in the job list.
5. Wait a for the Data Retention job to finish running.
6. From *Cluster Manager > Services > Data Lake*, manually *Restart* the Kudu service.
7. Check that the Kudu service has a healthy status.
8. If you still receive messages about the disk being full in the log, you might need to repeat steps 4-6.
9. When you stop receiving messages, go to *Cluster Manager > Services* and manually *Start* the Core service group.

How to fix Kudu metadata corruption

In rare situations, the Kudu tablet consensus may break and the metadata may corrupt, such as when an ungraceful host is powered off. When this occurs, *Monitor > Health > Database Health Check* will fail and report table conflict, unavailable, or under-replicated in the check result.

Example 1:

```
Tablet fcdb22e988f54674bf7bd81957d96d99 of table db_log_public.__root_fgt_hyperscale is conflicted:
  1 replicas' active configurations disagree with the leader master's:
  69c7e95e57f748bd801be9562db9684e (blade-10-0-1-6:7050): RUNNING
  538735a93bb8421b8fc2794fb31c52a7 (blade-10-0-1-5:7050): RUNNING
  077b5c932f2e4266820d13fb23442964 (blade-10-0-1-7:7050): RUNNING [LEADER]
All reported replicas are:
A = 69c7e95e57f748bd801be9562db9684e
B = 538735a93bb8421b8fc2794fb31c52a7
C = 077b5c932f2e4266820d13fb23442964
D = d06a84c881704e5d9f363a77cfd721d5
```

Example 2:

```
Tablet 13762b6c83fa4e18843bdabccb4ecdc6 of table 'db_log_public.__kavsprwq_fgt_traffic' is
unavailable: 2 replica(s) not RUNNING
2504215476754ff59ad3f0fcb0d58355 (blade-198-18-1-9:7050): RUNNING
04735952e86f4a4a98f748d1c2546fd8 (blade-198-18-1-8:7050): not running [LEADER]
State: FAILED
Data state: TABLET_DATA_READY
Last status: Corruption: Failed log replay. Reason: ...
```

```
799029fbd6e54dbf8d7c52f4bb837111 (blade-198-18-1-2:7050): not running
State: FAILED
Data state: TABLET_DATA_READY
Last status: Corruption: Failed log replay. Reason: ...
```

All reported replicas are:

```
A = 2504215476754ff59ad3f0fcb0d58355
B = 04735952e86f4a4a98f748d1c2546fd8
C = 799029fbd6e54dbf8d7c52f4bb837111
```

Example 3:

Tablet 2aaf23af4cc241e1865d292462418046 of table 'db_log_public.__kavsprwq_fgt_traffic' is **under-replicated**: 1 replica(s) not RUNNING

```
4df0f22bb51948fbbc0540383b0cfe74 (blade-198-18-1-5:7050): RUNNING
e84166ee4a1f4fb58753d5570273eaa1 (blade-198-18-1-10:7050): RUNNING [LEADER]
f02719025c0a41db92c64e0f37b94c53 (blade-198-18-1-4:7050): not running
State: FAILED
Data state: TABLET_DATA_READY
Last status: Corruption: Failed log replay. Reason: ...
```

All reported replicas are:

```
A = 4df0f22bb51948fbbc0540383b0cfe74
B = e84166ee4a1f4fb58753d5570273eaa1
C = f02719025c0a41db92c64e0f37b94c53
```

To fix the metadata corruption:

1. Rerun the Database Health Check multiple times and see if the issue persists. Sometimes Kudu may be able to heal by itself.
2. If the issue persists after a few retries, go to *Cluster Manager > Jobs > Create Custom Job*.
3. From the *Template* dropdown, select *Kudu Recover Corrupted Tablets* and Create.
4. In the *Jobs* table view, find the *Kudu Recover Corrupted Tablets* row, and click *Run in Actions*.
After the job is submitted, the tablet goes into recovering mode (see the example Database Health Check result below). The recovery may take several minutes, depending on the tablet size. Run the Database Health Check repeatedly until the health check returns success.

Example:

```
Tablet fcdb22e988f54674bf7bd81957d96d99 of table 'db_log_public.__root_fgt_hyperscale' is
recovering: 1 on-going tablet copies
69c7e95e57f748bd801be9562db9684e (blade-10-0-1-6:7050): not running
State: INITIALIZED
Data state: TABLET_DATA_COPYING
Last status: Tablet Copy: Downloading block 0000000022163966 (8961/25704)
538735a93bb8421b8fc2794fb31c52a7 (blade-10-0-1-5:7050): RUNNING
077b5c932f2e4266820d13fb23442964 (blade-10-0-1-7:7050): RUNNING [LEADER]
```



If all replicas fail, you cannot fix the metadata corruption. Instead, the entire tablet must be deleted in the recover job. See [To delete the entire tablet: on page 127](#).

In the below example, all replicas have failed.

```
Tablet 92ff614879504ee78b7c185a7d500d87 of table 'db_log_public.__8nhzlsno_facet_result' is
unavailable: 3 replica(s) not RUNNING
```

```
fadb45339c9e426b95f060d831faa4cc (blade-198-18-1-8:7050): not running
  State: STOPPED
  Data state: TABLET_DATA_TOMBSTONED
  Last status: Deleted tablet blocks from disk
780c0e107ce54dc384f75766906f4177 (blade-198-18-2-14:7050): not running [LEADER]
  State: STOPPED
  Data state: TABLET_DATA_TOMBSTONED
  Last status: Deleted tablet blocks from disk
bc3d03488e5a47909135510210f899f5 (blade-198-18-1-10:7050): not running
  State: STOPPED
  Data state: TABLET_DATA_TOMBSTONED
  Last status: Deleted tablet blocks from disk
```

To delete the entire tablet:

1. Go to *Cluster Manager > Jobs > Create Custom Job*.
2. From the *Template* dropdown, select *Kudu Recover Corrupted Tablets*.
3. Enable *Delete Entire Tablet*.
4. Click *Create*.
5. In the *Jobs* table view, find the *Kudu Recover Corrupted Tablets* row, and click *Run in Actions*.

How to enable/disable PXE boot server on Security Event Manager Controller

Security Event Manager Controller requires a PXE boot server to boot the network hosts in order to process deployment, scale-up or blade replacement. In a FortiAnalyzer-BigData 4500F and FortiAnalyzer-BigData 4500G appliance, this PXE boot server is disabled by default. You may want to enable it to allow certain operations, or disable it to avoid any unintended network boot in the FortiAnalyzer-BigData deployment’s internal network. To minimize unexpected side effect, PXE can only be enabled to run once. After the operation, PXE will be disabled again automatically.

To disable the PXE boot server:

Access the Security Event Manager Controller (see [To connect to the Security Event Manager Controller](#)) and run the following command to disable installation of system software:

```
fazbdadm disable pxe
```

To enable the PXE boot server to run once only:

Access the Security Event Manager Controller (see [To connect to the Security Event Manager Controller](#)) and run the following command to enable installation of system software:

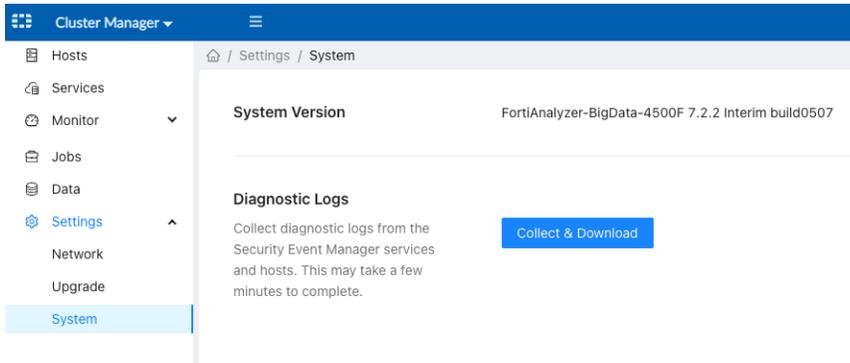
```
fazbdadm enable pxe-once
```

How to collect Diagnostic Logs

Diagnostic logs of the Security Event Manager services and host OSs can be collected from the Cluster Manager GUI.

To collect diagnostic logs via the GUI:

1. Go to *Cluster Manager > Settings > System*.
2. In the *Diagnostic Logs* section, click *Collect & Download*.



To collect diagnostic logs via the CLI:

1. Access the Security Event Manager Controller (see [To connect to the Security Event Manager Controller: on page 12](#)).
2. Enter the following command to collect the logs. Follow the output to locate the generated file with the logs packaged.

```
fazbdadm log export all
```

FAQ

Why does FAZBD4500F have physical ports 10G, 40G on SW module #1 & #2 when module #1 is for cluster internal connection?

By default, module #1 is for internal connections. If you need to expand the chassis, then you will need to link 2 boxes through the module #1 interfaces.

If there is no need for chassis expansion, you would not touch switch module #1.

Do port2 of FAZ Blade (slot 1) and Big Data hosts (Slots 2-14) have to be in the same subnet?

The recommended topology and the best practice is to use the same subnet.

If there is a second chassis (as for storage) are Blade 1 and Blade 2 used the same as Blade 3-14 or do they have special functions?

For the second chassis:

- Blade 1 is not used and should be powered off.
- Blade 2 does not have special functions, it is the same as Blades 3-14.

How many power supplies can fail without issues?

It is based on the workloads' power consumption. The field needs to observe the symptom and ensure the power supply to make it work.

Can we replace Blades/Switches for RMA?

Yes.

What happens if Blade 1 malfunctions?

If Blade 1 fails, it currently cannot failover to the other standby blade. However, as of v7.0.1, if you stack two FAZ-BD chassis, HA mode can be enabled across the two FAZ blades.

If Blade 2 fails, does it delegate its role to another blade?

If Blade 2 fails, it can failover to the other Big Data host.

Big Data controller's IP is configured in FortiAnalyzer Blade 1. Does this mean that this IP is "shared" between all Big Data hosts (Blades 2-14)?

The controller IP is not shared. It is configured on the Big Data controller. When the controller fails over, the IP is re-configured on a new controller.

When a blade fails, does it delegate its role to another blade?

When Blades 2-14 fails, the node with the same role will take over.

Does "3 K8s Master" mean there are 3 blades for this function?

Yes, it means there are 3 blades/hosts that are masters.

Is "3" the default replication factor?

Yes, there are 3 replications.

This means that 3 copies are stored in the system:

- 1 original copy
- 2 replicated copies

Is the allocated disk usage for the FortiAnalyzer blade or for the Big Data blade?

It is for FortiAnalyzer to cache the log data before they are ingested to the Big Data cluster.

To view the allocated disk usage, go to *Settings > Manage Storage Policy > Data Policy > Disk Usage > Allocated*.

In Hosts view, is the last number the actual slot?

The last number is the blade ID. It is important to install blades into slots according to IDs.

Is the blade ID labeled on the blades?

Yes.

If the blade is labeled, when a particular faulty blade is on RMA, should the RMA return with the same blade labeled with the same blade ID?

The RMA will ship a blade with the OS and bootloader wiped, so all the Big Data blades are the same no matter which blade ID it belongs to. Once it is injected into the slots, you will need to set the blade ID accordingly in bootloader. The Big Data controller will bootstrap it with a cluster role.

Does this mean HA is no longer supported?

For 6.2 & 6.4, the Blade 2-14 HA is supported by the Big Data architecture design.



Previous to v6.4.6, recovering/replacing a non-Data role node may fail and require a manual recovery process in some scenarios. This has been resolved in v6.4.6.

Can I migrate old logs after removing fetch management?

Data migration is not supported.

Is the upgrade started from the Big Data controller going to upgrade all the blades including Blade 1?

Yes. CLI runs on the controller node, then it will upgrade all the nodes, starting with and including Blade 1.

Will the size of the database affect the amount of time it takes to upgrade? For example, will it take less time to upgrade a 1Tb vs 150Tb?

Based on the Big Data architecture, the database is not rebuilt. Therefore, the size of the database will not impact the upgrade time.

Will Log View and Reports show data from its own cache and Big Data storage?

Data is stored in Big Data storage. When you query the log data from *Log View* and *Reports*, the query request will be sent to the Big Data. All the data returned comes from the Big Data cluster.

How much free storage is required for upgrade?

The system has a mechanism to reserve space for upgrade and other operations.

For upgrade, will all Big Data blades reboot at the same time, or in the order of the blade that finished downloading the image first?

The infrared controller handles the order of the blades. They will be rebooted parallel of each other.

When the system is working normally (not specific to upgrade), is log ingestion from the ADOM cache to Big Data storage in real time or is there a delay?

It is near real time. The delay will be several minutes at most.

Does "same version" mean same main branch (v7.4 to v7.4) or same patch level (v7.4.1 to v7.4.1)?

"Same version" means same patch level GA release.

Will the system rollback if an upgrade fails for certain blades?

The system will not rollback. However, you can use the CLI console to redo upgrading to recover.

```
fazbdctl upgrade fazbd -o retry
```

Change Log

Date	Change Description
2024-07-02	Initial release v7.4.1.
2024-07-10	Updated Built-in automation jobs on page 69 .
2024-10-16	Updated Alert on page 37 .
2025-01-20	Removed reference to FortiAnalyzer-BigData-VM.
2025-02-26	Updated FortiAnalyzer-BigData hardware environment on page 11 .
2025-03-27	Updated How to fix Kudu metadata corruption on page 125 .
2025-05-05	Updated How to enable/disable PXE boot server on Security Event Manager Controller on page 127 .
2025-08-08	Updated How to scale out on page 103 .



www.fortinet.com

Copyright© 2025 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.