

Administration Guide

FortiAnalyzer-BigData 7.0.1



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



January 23, 2024

FortiAnalyzer-BigData 7.0.1 Administration Guide

58-701-744036-20240123

TABLE OF CONTENTS

About FortiAnalyzer-BigData	6
Main Features	6
Supported models	7
Key terms and concepts	7
FortiAnalyzer-BigData Hardware environment	12
Set up process	13
Initial set up	13
Set up the FortiAnalyzer-BigData network	14
Set up Administrator accounts	15
Connect to the Chassis Management Module	16
Set up the CMM network	16
Configure the CMM password	20
Configure the Blade Management Network	21
Remotely control blades via CMM	22
Configure the BMC password	23
Turn off STP BPDU	27
Configure LACP	28
Connect to the FortiAnalyzer-BigData CLI	31
GUI overview	32
Cluster Manager	32
Custom refresh settings	34
Commands management	34
Notifications management	35
Host management	36
Role assignment	37
Service management	38
Service groups	39
Service details	40
Monitor	43
Dashboards	43
Filtering the Dashboard	44
Customizing the Dashboard	44
Logs and metrics	46
Explore logs	47
Explore metrics	51
Health	53
Health Check	53
Alert	54
Hyperscale firewall NetFlow logging support	59
Set up Security Manager hosts external IP addresses	59
Configure FortiAnalyzer-BigData as IPFIX log server on FortiGate	61
Device Manager and log rate	62

Search IPFIX log in Log View	63
Global search	64
Starting a global search	64
Global Search settings	64
Log types (Global Search)	65
Histogram	65
Event Inspector	66
Faceted Search (+, -, focus)	66
Time Window	68
Search History	68
Split View	69
Live Streaming Search	70
Cross-Cluster Search Federation	70
Create a new Search Federation (Example)	71
Search with Federation	76
Log Query Language (LogQL)	77
Log Stream Selector	77
Log Pipeline	78
Job management and automation	80
Job history	81
Built-in automation jobs	82
Custom automation jobs	82
Custom job templates	84
Data management	88
Manage storage policy	90
Data backup	90
Incremental backups	92
Incremental backups	93
Data restore	96
Bootloader	99
Bootloader Main Page	99
1. Configure Network	100
2. Install OS	100
3. Set Role	101
4. Set Chassis ID	101
5. Set Blade ID	101
6. Reset OS	101
7. Reset OS and Clear User Data	102
8. Upgrade Bootloader	102
9. Reboot	102
sh. shell	102
General maintenance and best practices	103
Backup and restore to external HDFS	103
Schedule maintenance tasks for off-peak hours	103
Maintain database integrity	104

Upgrade FortiAnalyzer-BigData	105
Scaling FortiAnalyzer-BigData	108
How to scale out	108
How to remove a chassis from a stacked setup	109
Remove an extender chassis	109
Reset FortiAnalyzer-BigData	110
Soft reset FortiAnalyzer-BigData	110
Hard reset FortiAnalyzer-BigData	110
Troubleshooting	112
What to do if an upgrade fails	112
What to do if a soft reset fails	112
What to do if a hard reset fails	113
How to repair disk failures	114
How to replace a blade	114
How to reset a single host	115
How to rebalance the data	115
How to recover from an unhealthy service status	116
Core services	116
Data Lake services	117
Message Broker services	118
How to recover from a full disk	119
How to fix Kudu consensus mismatch	119
How to set up management and external IP addresses using CLI	120
Setting up management IP address on the Security Event Manager Controller	120
Setting up external IP address on a single Security Event Manager host	121
Setting up external IP addresses on all Security Event Manager hosts	121
Clearing external IP addresses on Security Event Manager hosts	122
Displaying external IP addresses on Security Event Manager Controller and hosts	122
Change Log	123

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.

Supported models

FortiAnalyzer-BigData supports the same FortiGate models as FortiAnalyzer 7.0.1. For a list of supported FortiGate models, see the [FortiAnalyzer 7.0.1 Release Notes](#).

Key terms and concepts

This section contains key terms used in FortiAnalyzer-BigData.


Security Event Manager

The Security Event Manager is formed by Blade A2–A14 that serves 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 is a single host within the Security Event Manager that functions as the main controller for the hosts. This host is usually Blade A2 of the chassis and is responsible for the DHCP, configuration management, and lifecycle management such as upgrades, resets, and more. If this host goes down, it will automatically failover to a standby host.

When a failover occurs, the Controller can be one of the nodes other than Blade A2. 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 on any of the Security Event Manager hosts:

```
fazbdctl show members
```

The controller appears in the *Role* column

Security Event Manager Hosts

This refers to Blade A2–A14, which are the hosts that form the Security Event Manager.

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. 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 runs on Blade A1 and is responsible for collecting logs and providing the services for FortiView, Log View, Reports, and more.

Roles

The Security Event Manager hosts are categorized into three 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 108](#)), the hosts on the extender chassis can be added as data nodes to the existing cluster in the main chassis.

FortiAnalyzer-BigData has the following roles and 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


The FortiAnalyzer-BigData 4500F unit is a 4U chassis with two 10G network switch modules, and 14 blades in the enclosure.

Each blade contains two 2.1GHz Intel Xeon 8 Core 16 Thread (8C16T) CPU, 128GB RAM, and two 7.68TB NVMe SSD.

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

The two network switch modules 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 [Connect to the Chassis Management Module on page 16](#).

Set up process

The set up process for FortiAnalyzer-BigData consists of setting up the FortiAnalyzer-BigData unit and the [Chassis Management Module \(CMM\)](#).

To set up the FortiAnalyzer-BigData unit, you must perform the following steps:

1. [Initial set up on page 13](#)
2. [Set up the FortiAnalyzer-BigData network on page 14](#)
3. [Set up Administrator accounts on page 15](#)

Once the unit and network is set up and connected, you can [connect to the Main CLI or Security Event Manager Controller](#).

In addition to setting up FortiAnalyzer-BigData, you also need to [set up the Chassis Management Module \(CMM\)](#).

Prerequisites

You must have the following before beginning to set up your FortiAnalyzer-BigData:

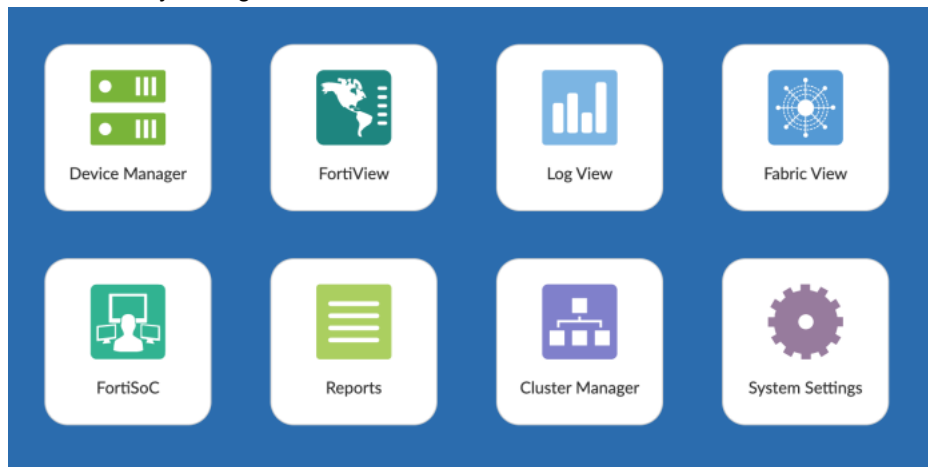
- Ethernet cable
- SPF RJ45 transceiver module
- Management computer

Initial set up

To connect to the FortiAnalyzer-BigData GUI:

1. Install the SFP RJ45 transceiver module into one of the SFP interfaces on the FortiAnalyzer-BigData Switch Module #2.
2. Connect the RJ45 port on the transceiver module to the management computer using the supplied Ethernet cable.
3. Enable DHCP or set the management computer's IP address to be on the same subnet as FortiAnalyzer-BigData.
For example:
 - **IP address:** 192.168.1.10
 - **Netmask:** 255.255.255.0
4. On the management computer, open a supported web browser and visit <https://192.168.1.98>.
5. Log in with the username `admin` and no password.

The FortiAnalyzer-BigData GUI loads.



Set up the FortiAnalyzer-BigData network

To set up the network for FortiAnalyzer-BigData, users need to connect either a 10GE link with SFP, or 40GE link with QSFP, from Switch Module #2 to your public access switch, and then set up the external IP address via the FortiAnalyzer-BigData GUI. This setup requires two IPs from the same subnet for logging (*Main Host*) and management (*Security Event Manager*) access.

To set up the FortiAnalyzer-BigData network:

1. From the FortiAnalyzer-BigData GUI, go to *System Settings > Network*.
2. Change the *Security Event Manager IP Address/Netmask* and *Gateway* fields to your internal network.
This is the address of the FortiAnalyzer-BigData Security Event Manager which is responsible for serving the web GUI and performs various data processing and management workload.
3. Change the *Main Host IP Address/Netmask* and *Gateway* fields to your internal network.
This is the address of the FortiAnalyzer-BigData Main host, which is responsible for collecting the log and serving the GUI for FortiView, LogView, Reports, and so on.
4. Keep the default *Administrative Access* settings.
5. Specify a *Default Gateway and DNS Servers*.
6. Click *Apply* to save your changes.

7. From your management computer, change the IP Address/Netmask to reconnect it to FortiAnalyzer-BigData.

System Network Management Interface

Name	port2
Security Event Manager IP Address/Netmask	<input type="text" value="10.105.101.59/255.255.255.0"/>
Security Event Manager Default Gateway	<input type="text" value="10.105.101.1"/>
Main Host IP Address/Netmask	<input type="text" value="10.105.101.58/255.255.255.0"/>
Main Host Default Gateway	<input type="text" value="10.105.101.1"/>
IPv6 Address	<input "::="" 0"="" type="text" value=""/>
Administrative Access	<input type="checkbox"/> HTTPS <input type="checkbox"/> HTTP <input checked="" type="checkbox"/> PING <input checked="" type="checkbox"/> SSH <input checked="" type="checkbox"/> SNMP <input type="checkbox"/> Web Service <input checked="" type="checkbox"/> FortiManager
IPv6 Administrative Access	<input type="checkbox"/> HTTPS <input type="checkbox"/> HTTP <input type="checkbox"/> PING <input type="checkbox"/> SSH <input type="checkbox"/> SNMP <input type="checkbox"/> Web Service <input type="checkbox"/> FortiManager
Primary DNS Server	<input type="text" value="208.91.112.52"/>
Secondary DNS Server	<input type="text" value="208.91.112.53"/>

All Interfaces

Routing Table

IPv6 Routing Table

Packet Capture

Apply

Set up Administrator accounts

Set up an administrator account so you can configure your FortiAnalyzer-BigData.

To set up an Administrator account:

1. Go to *System Settings > Admin > Administrators*, and click *Create New* in the toolbar.
2. In the *User Name* field, enter a new name for your administrator.

3. In the *New Password* and *Confirm Password* fields, enter the password for the administrator account.

New Administrator

User Name

Example_Admin

Avatar

E

+ Change Photo

- Remove Photo

Comments

0/127

Admin Type

LOCAL

New Password

••••••••

Confirm Password

••••••••

Admin Profile

Restricted_User

Administrative Domain

All ADOMs

All ADOMs except specified ones

Specify

Trusted Hosts

OFF

Meta Fields >

Advanced Options >

OK

Cancel

4. Click **OK** to save.

Connect to 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. The CMM comes with a web management utility that consolidates and simplifies system management for the FortiAnalyzer-BigData chassis. This setup requires 15 IPs from the same subnet: One CMM IP and 14 Blade IPMI IPs, and an additional two IPs for two switch modules' management GUI if needed.

Set up the CMM network

To set up CMM network via GUI:

1. Connect a 10GE link from the CMM module (the module in the middle of the back panel) to your public access switch, and set up the external IP address via the CMM web management utility.
2. Connect the port on the CMM Module to a management computer using the supplied Ethernet cable
3. Set the management computer's IP and subnet to be on the same subnet as FortiAnalyzer-BigData:
For example:
 - **Static IP Address:** 192.168.1.x
 - **Subnet Mask:** 255.255.255.0
4. On the management computer, open a supported web browser and visit <https://192.168.100.100> (the default CMM IP).

5. Log in with the default username and password on the Fortinet Product Credentials card.



Changing the default password is strongly recommended. See [Configure the CMM password on page 20](#).

6. Go to *Configuration > CMM Network* to configure the CMM network.
7. Select a radio button option for how you want to obtain at IP address.

CMM Network

This page you can view and modify the network settings. Select whether to obtain an IP address automatically or manually configure one.

MAC Address

Hostname

- ☐ Obtain an IP address automatically (use DHCP mode)
- ☐ Use the following IP address (use Static mode)
- ☒ Use the following IP address when DHCP fails(use Static mode when DHCP fails)

- **Obtain an IP address automatically:** Uses DHCP to automatically obtain the IP address.
- **Use the following IP address:** Set up the IP address by manually entering the IP information into the fields below.
- **Use the following IP address when DHCP fails:** If CMM is unable to obtain the dynamic IP from the DHCP server, it will use the static IP instead. This is the default setting.

8. Depending on the option you selected in step 6, enter your IP information under *IPv4 Setting*, *IPv4 Setting when DHCP fails*, or *IPv6 Setting*.

CMM Network

This page you can view and modify the network settings. Select whether to obtain an IP address automatically or manually configure one.

MAC Address

Hostname

☐ Obtain an IP address automatically (use DHCP mode)
☐ Use the following IP address (use Static mode)
☒ Use the following IP address when DHCP fails (use Static mode when DHCP fails)

IPv4 Setting

IP Address
 Subnet Mask
 Gateway
 DNS Server IP

IPv4 Setting when DHCP fails

IP Address
 Subnet Mask
 Gateway

IPv6 Setting

IPv6 Address
☒ Add IP ☐ Delete IP ☒ Auto Configuration
☒ DHCPv6 Stateless ☐ DHCPv6 Stateful
 Address List
 DNS Server IP
 DUID

VLAN ☐ enable ☒ disable

VLAN ID

RMCP Port

Network Link Status

Active Interface Dedicated
 Status : Connected
 Speed : 1G
 Duplex : Full Duplex

[Save](#)

9. If you need Virtual LAN support, select **enable** to enable VLAN and enter the VLAN ID in the field.

VLAN ☐ enable ☒ disable

VLAN ID

10. In the RMCP Port field, enter the desired Remote Mail Checking Protocol (RMCP) port based on your configuration. The default port is 623.
11. Once you are done completing the fields, click **Save** to save the CMM Network settings.

To set up CMM network via CLI:

1. Using a USB-to-RJ45 serial adapter, connect a management computer to the serial port on the CMM module.
2. Establish a serial connection to the CMM from the management computer using a serial terminal such as Putty or Hyper Terminal, and enter the following configuration.

Configure the serial line

Speed (baud)	<input type="text" value="115200"/>
Data bits	<input type="text" value="8"/>
Stop bits	<input type="text" value="1"/>
Parity	<input type="text" value="None"/>
Flow control	<input type="text" value="XON/XOFF"/>

3. Using the CMM CLI commands, set up IP addresses on the management port.

Example settings:

```
SET IP 10.100.100.099
SET NETMASK 255.255.255.0
SET GATEWAY 10.100.100.1
SET DHCP DISABLE
APPLY SETTING
```

CMM CLI Commands	Description
HELP	Print help.
RESET	Reset CMM.
DEFAULTRESET	Reset CMM to default.
VER	Show CMM FW VER.
PASSWORDRESET	Reset password.
GET LAN INFO	Get network info.
SET IP xxx.xxx.xxx.xxx	Set IP address.
SET NETMASK xxx.xxx.xxx.xxx	Set netmask address.
SET GATEWAY xxx.xxx.xxx.xxx	Set gateway address.
SET MAC xx:xx:xx:xx:xx:xx	Set MAC address.
SET DHCP ENABLE	Set DHCP enable.
SET DHCP DISABLE	Set DHCP disable.
SET DHCP FAILOVER	Set DHCP fails, then use manual configuration.
APPLY SETTING	Apply network setting.

4. Verify the network setup with the `GET LAN INFO` command.
5. Verify that the web management utility can be accessed from a web browser.

Configure the CMM password

You can configure the CMM password via the GUI or CLI.

To change the CMM password via GUI:

1. From a web browser, access the web management utility using the CMM IP address.
2. Log in with the admin username and password.
3. Go to *Configuration > Users*.

➔ Users

This page displays the list below shows the current list of configured users. If you would like to delete or modify a user, select their name in the list and press **[Delete User]** or **[Modify User]**. To add a new user, select an unconfigured slot and press **[Add User]**.

User ID	User Name	Network Privilege	Email
1	Anonymous	Reserved	~
2	ADMIN	Administrator	~
3	~	Reserved	~
4	~	Reserved	~
5	~	Reserved	~
6	~	Reserved	~
7	~	Reserved	~
8	~	Reserved	~
9	~	Reserved	~
10	~	Reserved	~

Number of Configured Users: 10

Add User **Modify User** **Delete User**

4. Select the *ADMIN* row and click *Modify User*.
5. Click the *Change Password* checkbox, change the password, and click *Modify*.

➔ Modify User

Enter the new information for the user below and press **[Modify]**. Press **[Cancel]** to return to the user list.

User Name:

Change Password ☒

Password:

Confirm Password:

Network Privileges:

Email Address:

Modify **Cancel**

To reset the CMM password via CLI:

1. Using a USB-to-RJ45 serial adapter, connect a management computer to the serial port on the CMM module.
2. Establish a serial connection to the CMM from the management computer using a serial terminal such as Putty or Hyper Terminal.
3. Use the `PASSWORDRESET` command to reset the password to the default password.

Configure the Blade Management Network



The Blade Management Network should be in the same subnet as Chassis Management Network. See, [Connect to the Chassis Management Module on page 16](#)

To configure the Blade Management Network for a single blade:

1. From a web browser, access the web management utility using the CMM IP address.
2. Go to *Blade System* > *Blade Status*, and click any blade you would like to set up the Blade Management network.
3. In the Blade Configuration view, click *Network Config*.
4. Configure the IPv4 Setting:
 - a. Enter the *IP Address*, *Subnet Mask*, *Gateway* and *DNS Service IP*.
 - b. Click *Save*.

Hide >>> [Blade A1 Node --- 1] Summary Sensor Reading Network Config Health Event Log Maintenance Event Log FRU Info HW Information Reset Default Configuration

➔ Blade IPMI Network

This page you can view and modify the network settings. Select whether to obtain an IP address automatically or manually configure one.

MAC Address

Hostname

☐ Obtain an IP address automatically (use DHCP mode)
☒ Use the following IP address (use Static mode)

IPv4 Setting

IP Address	<input type="text" value="010.100.100.100"/>
Subnet Mask	<input type="text" value="255.255.255.000"/>
Gateway	<input type="text" value="010.100.100.001"/>
DNS Server IP	<input type="text" value="172.001.001.100"/>

To configure the Blade Management Network for all blades at once:

1. From a web browser, access the web management utility using the CMM IP address.
2. Go to *Configuration* > *Blade IPMI Network* to access the Blade IPMI Network page.

Blade IPMI Network

This page you can configure Blade IPMI network settings.

☒ Use the following IP address (use Static mode)

IPv4 Setting

IP Scale	<input type="text" value="1"/>
Base IP Address	<input type="text" value="010.100.100.100"/>
Subnet Mask	<input type="text" value="255.255.255.000"/>
Gateway	<input type="text" value="010.100.100.001"/>
DNS Server IP	<input type="text" value="172.001.001.100"/>
VLAN ID	<input type="text" value="0"/>

Save

The Blade IPMI Network page enables you to modify the Blade Management Controller (BMC) networks of all your blades.

3. Configure the IPV4 settings.
 - a. Select *Use the following IP address*.
 - b. From the *IP Scale* dropdown, select the base number each blade IP address will increase (1, 2, or 4).
 - c. Enter the *Base IP Address*, *Subnet Mask*, *Gateway*, *DNS Service IP*, and *VLAN ID*.
The *Base IP Address* is applied to the first node of a blade's A1 and increases by a set amount for every following node.
4. Click **Save** and accept the warning prompts.

Remotely control blades via CMM

The CMM web management utility can perform various remote operations on the chassis, such as remote console and power control. This can be used for running diagnostic tasks on individual blades. It also allows the administrator to

remotely control the FortiAnalyzer-BigData via CLIs if the Main IP and the BigData Controller IP are reset after a software hard reset.

To access the FortiAnalyzer-BigData Main CLI:

1. Go to *Blade System > Summary* and select *Blade A1*.
2. To enter the BMC for the FortiAnalyzer-BigData Main Host, click the *BMC IPV4* link.
3. Enter your username and password to log in.
The default login credentials are on the Fortinet Product Credentials card.
4. Go to *Remote Control > Console Redirection or iKVM/HTML5*.
5. Log in with username `admin` and no password.
You can now configure the Main host via the CLI.

To access the Security Event Manager Controller:

1. Go to *Blade System > Summary* and select *Blade A2*.
2. To enter the BMC for the Security Event Manager Controller, click the *BMC IPV4* link.
The default login credentials are on the Fortinet Product Credentials card.
3. Go to *Remote Control > Console Redirection or iKVM/HTML5*.
4. Log in with username `root` and password `fortinet@123`.
You can now access the Security Event Manager Controller and use `fazbdctl` CLI commands to manage the cluster.



You can use the CMM web management utility to remotely access and control the other blades by following the general steps.

You can also use the utility to remotely access the FortiAnalyzer-BigData Bootloader (see [Bootloader on page 99](#)).

Configure the BMC password

You can configure the BMC password via the CMM.

To change the BMC password via the CMM:

1. From a web browser, access the web management utility using the CMM IP address.
2. Log in with the admin username and password.
3. Go to *Blade System > Summary*.
4. Select the blade you want to change, for example, Blade A1.
5. To enter the BMC for the FortiAnalyzer-BigData main host, click the *BMC IPV4* link.
The default login credentials are on the Fortinet Product Credentials card.

6. Go to *Configuration > Users*.

➔ Users

This page displays the list below shows the current list of configured users. If you would like to delete or modify a user, select their name in the list and press **[Delete User]** or **[Modify User]**. To add a new user, select an unconfigured slot and press **[Add User]**.

Number of Configured Users: 10

User ID ↕	User Name ↕	Network Privilege ↕
1	Anonymous	Reserved
2	ADMIN	Administrator
3	~	Reserved
4	~	Reserved
5	~	Reserved
6	~	Reserved
7	~	Reserved
8	~	Reserved
9	~	Reserved
10	~	Reserved

Add User **Modify User** **Delete User**

7. Select the *ADMIN* row and click *Modify User*.8. Click the *Change Password* checkbox, change the password, and click *Modify*.

➔ Modify User

Enter the new information for the user below and press **[Modify]**. Press **[Cancel]** to return to the user list.

User Name:

Change Password ☒

Password:

Confirm Password:

Network Privileges:

Modify **Cancel**

To reset the BMC password via CMM:

1. From a web browser, access the web management utility using the CMM IP address.
2. Log in with the admin username and password.

3. Go to *Blade Status* and select the blade you want to change, for example, Blade A1.

Blade Status

Power Off | Power On | Power Cycle | Power Reset | Graceful Shutdown | AC Cycle | PwrFail Policy | Pwr Capping | ACLost Policy | Refresh | Auto Refresh

Blade	Name	Model	Pwr Status	Max Pwr	iKVM/HTML5	UID	Status	BMC IPv4	BMC IPv6	BMC Ver
<input type="checkbox"/> Blade A1		FAZ-BD	On/Off	289	iKVM HTML5	Off UID	Normal	10.105.101.101	fe80::ae1f:6bff:fec0:bc1a/64	13.72.00
<input type="checkbox"/> Blade A2		FAZ-BD	On/Off	289	iKVM HTML5	Off UID	Normal	10.105.101.102	fe80::ae1f:6bff:fec0:ba78/64	13.72.00
<input type="checkbox"/> Blade A3		FAZ-BD	On/Off	289	iKVM HTML5	Off UID	Normal	10.105.101.103	fe80::ae1f:6bff:fec0:ba63/64	13.72.00
<input type="checkbox"/> Blade A4		FAZ-BD	On/Off	289	iKVM HTML5	Off UID	Normal	10.105.101.104	fe80::ae1f:6bff:fec0:ba6c/64	13.72.00
<input type="checkbox"/> Blade A5		FAZ-BD	On/Off	289	iKVM HTML5	Off UID	Normal	10.105.101.105	fe80::ae1f:6bff:fec0:ba6e/64	13.72.00
<input type="checkbox"/> Blade A6		FAZ-BD	On/Off	289	iKVM HTML5	Off UID	Normal	10.105.101.106	fe80::ae1f:6bff:fec0:ba36/64	13.72.00
<input type="checkbox"/> Blade A7		FAZ-BD	On/Off	289	iKVM HTML5	Off UID	Normal	10.105.101.107	fe80::ae1f:6bff:fec0:ba6d/64	13.72.00
<input type="checkbox"/> Blade A8		FAZ-BD	On/Off	289	iKVM HTML5	Off UID	Normal	10.105.101.108	fe80::ae1f:6bff:fec0:ba96/64	13.72.00
<input type="checkbox"/> Blade A9		FAZ-BD	On/Off	289	iKVM HTML5	Off UID	Normal	10.105.101.109	fe80::ae1f:6bff:fec0:bac6/64	13.72.00
<input type="checkbox"/> Blade A10		FAZ-BD	On/Off	289	iKVM HTML5	Off UID	Normal	10.105.101.110	fe80::ae1f:6bff:fec0:ba91/64	13.72.00
<input type="checkbox"/> Blade A11		FAZ-BD	On/Off	289	iKVM HTML5	Off UID	Normal	10.105.101.111	fe80::ae1f:6bff:fec0:ba9c/64	13.72.00
<input type="checkbox"/> Blade A12		FAZ-BD	On/Off	289	iKVM HTML5	Off UID	Normal	10.105.101.112	fe80::ae1f:6bff:fec0:bae0/64	13.72.00
<input type="checkbox"/> Blade A13		FAZ-BD	On/Off	289	iKVM HTML5	Off UID	Normal	10.105.101.113	fe80::ae1f:6bff:fec0:ba26/64	13.72.00
<input type="checkbox"/> Blade A14		FAZ-BD	On/Off	289	iKVM HTML5	Off UID	Normal	10.105.101.114	fe80::ae1f:6bff:fec0:ba77/64	13.72.00

4. Click *Reset Default Configuration*.

Hide >>> [Blade A1 Node 1] --- Summary | Sensor Reading | Network Config | Health Event Log | Maintenance Event Log | FRU Information | Date & Time | Power/Temp Record | Node Product Key

HW Information | **Reset Default Configuration**

Node Status and Control

Location:	A1-1		
Board Model:	4500PT		
Product Model:	FAZ-BD		
Blade Max Pwr:	289		
Blade Curr Power:	108		
Error:	Normal		
Post Code:	FF		
BMC Version:	13.72.00		
CPLD Version:	02.b3.05		
BMC IPv4 Addr:	10.105.101.101	BMC Reset BMC Reset To Default	
BMC IPv6 Addr:	fe80::ae1f:6bff:fec0:bc1a/64		
KVM:	Not Launched	KVM Launch	
VM:		VM Launch	
SOL:		SOL Launch	
Blade UID:	Off	UID Off UID On	
Node Name:		Save Node Name	
Blade Name:		Save Blade Name	
PwrFail Policy:	Throttle	Save PwrFail Policy	
Pwr Status:	On	Power Off Power On Reset Power Cycle Graceful Shutdown	

Motherboard Information

BIOS		CPU		Memory		Onboard NIC	
BIOS ID	4500PT	Num of CPU	2	Num of DIMM	8	Num of NIC	2
BIOS Version	3.3	CPU ID	0654	Memory Size	131072 MB	NIC1 MAC	ac:1f:6b:5a:a1:28
Build Date	06/01/2020	CPU Speed	2100 Mhz	Memory Speed	2666 Mhz	NIC2 MAC	ac:1f:6b:5a:a1:29
						NIC3 MAC	N/A
						NIC4 MAC	N/A

Refresh | Auto Refresh

5. Select the *Reset Users Configuration* checkbox and click *Reset*.

Hide
>>>

**[Blade A1
Node 1]**

Summary

Sensor Reading

Network Config

Health Event Log

-

HW Information

Reset Default Configuration

➔ Reset Default Configuration

This page is to reset blade configuration to defaults settings by clicking on the [Reset] button.

- ☐ Reset All Configurations below
- ☐ Clear Power/Temperature Record Clear peak record ▼
- ☐ Reset Health Event Log and Configuration
- ☐ Reset Maintenance Event Log and Configuration
- ☐ Reset Alert Configuration
- ☐ Reset Date&Time Configuration
- ☐ Reset LDAP Configuration
- ☐ Reset Active Directory Configuration
- ☐ Reset RADIUS Configuration
- ☐ Reset Mouse mode Configuration
- ☐ Reset Network Configuration
- ☐ Reset Dynamic DNS Configuration
- ☐ Reset SMTP Configuration
- ☒ Reset Users Configuration
- ☐ Reset Port Configuration
- ☐ Reset IP Access Control Configuration
- ☐ Reset SNMP Configuration
- ☐ Reset Web Session Configuration
- ☐ Reset SDR Configuration
- ☐ Clear SSL Certification Configuration
- ☐ Reset RAKP Configuration
- ☐ Reset HTTPD Configuration
- ☐ Reset Syslog Configuration

Reset

6.

Turn off STP BPDU

To turn off STP BPDU:

1. Connect to the Chassis Management Module on page 16
2. Go to *Blade System > Switch Module*, and click *Switch A2*. The *Switch Module* pane opens.

Fortinet

CMM Name: 10.105.101.80
User: ADMIN (Administrator)

Blade System > System Health > Configuration > Remote Control > Maintenance > Help

Switch Module

HW Reset UID On UID Off Refresh Auto Refresh

Switch

Switch	Switch Type	Module
Switch A1	10G Ethernet Switch	MBM-XE
Switch A2	10G Ethernet Switch	MBM-XE

Hide >>> [Switch A1] Summary FRU Information

Switch Module

Switch Information

Switch	Switch Type	Module Name	Pwr Status	Temperature	UID	Error	Management IP	FW Ver	Pwr Cons
[Switch A1]	10G Ethernet Switch	MBM-XEM-002	On	45.39	Off	Normal	10.105.101.75	2.1.0-73	46 W

HW Reset UID On UID Off Refresh Auto Refresh

Configure Date and Time Settings.

☐ CMM

☒ Local

Date and Time: 01-28-2000 22:37:31

☐ Apply above setting to all Switches

Save

Switch Network Configuration

IP Config: Use the following IP address (use Static mode)

IP Address: 10.105.101.75

Subnet Mask: 255.255.255.0

Gateway: 10.105.101.1

Mgmt 1 MAC Address: ac:1f:6b:f0:b7:58 (Activated)

Mgmt 2 MAC Address: ac:1f:6b:f0:b7:58

Save

Switch Username & Password Reset

Username: ADMIN

Password: [password field]

Confirm Password: [password field]

Save

Reset to Factory Default



The default *Username* and *Password* are both ADMIN.
For security purposes, we recommend changing the *Username* and *Password*.

3. Under *Switch Network Configuration*, in the *IP Address* field, enter the IP address, and click *Save*.
4. Under *Switch Information*, click the *Management IP* column, and enter the management web GUI for *Switch A2*.
5. Go to *Layer-2 > MSTP > Basic Settings*.
 - a. Set *MSTP Status* to *Disabled*.
 - b. Set *System Control* to *Shutdown*.

c. Click *Apply*.

Select	Context Id	System Control	MSTP Status	Maximum MST Instances	Bridge Priority	Protocol Version	Region Name	Region Version	Dynamic Path Cost Calculation	Speed Change Path Cost Calculation
<input checked="" type="radio"/>	0	Shutdown	Disabled	64	32768	MSTP	ac:1f:6b:92:38:fa	0	True	True

Apply

Configure Trace Options

6. Go to *Layer-2 > RSTP > Global Settings*, and confirm:

- *Status is Disabled*
- *System Control is Shutdown* (default)

Select	Context Id	System Control	Status	Dynamic Path Cost Calculation	Speed Change Path Cost Calculation	Flush Interval	Flush Indication Threshold	BPDU Guard
<input checked="" type="radio"/>	0	Shutdown	Disabled	True	True	0	0	

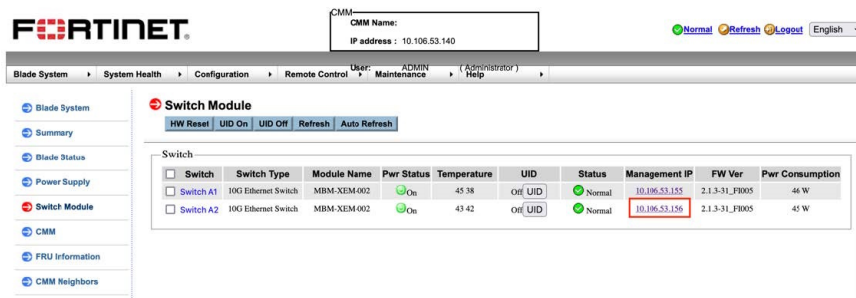
Apply

Configure Trace Options


Configure LACP

To configure port channel on the FortiAnalyzer-BigData switch module:

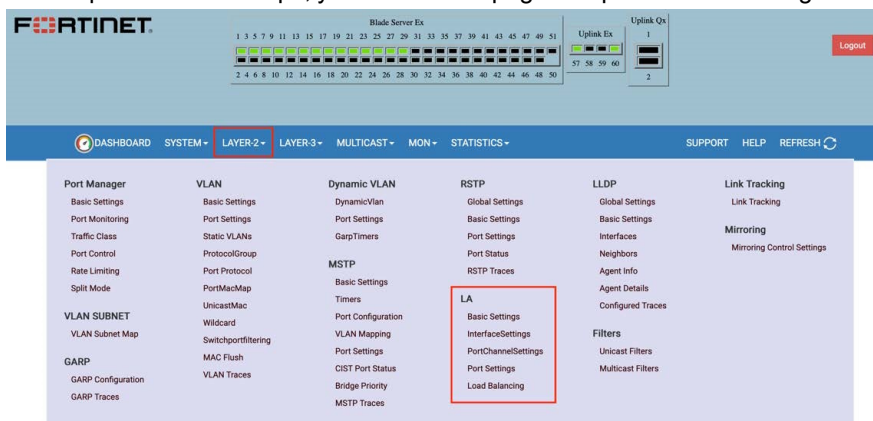
1. In CMM, go to *Switch Module* and click the Management IP of Switch A2 to log into the switch web-based management utility.



2. In the switch web-based management utility, the switch ports are displayed in home page. The switch external ports are Ex0/57 to Ex0/60, and Qx0/1 to Qx0/2. Mouse over the port to get the port name.



3. To complete the next steps, you will use the pages for port channel configurations under *LAYER-2 > LA*.



4. In *LA Basic Settings*, verify the following:

- *System Control* is *Start*
- *LA Status* is *Enabled*

LA BASIC SETTINGS

System Control	Start
LA Status	Enabled
System Priority	32768
System ID	ac:1f:6b:f3:2a:51
Apply	

5. In *LA Interface Settings*, input *Port Channel ID*, and click *Add*.

PORTCHANNEL INTERFACE BASIC SETTINGS

Port Channel ID	100 *
Context ID	0 ▾
Admin Status	Up ▾
MTU	
<div>Add Reset</div>	

6. In *LA Port Channel Settings*, input member ports in *Ports*, and click *Apply*.

LA PORT CHANNEL SETTINGS

Port Channel ID	po100 ▾ *
Aggregation Type	Static ▾
Action Type	Add ▾
Mode	Lacp ▾
Ports	Qx0/1,Qx0/2
DefaultPort	Qx0/2 ▾
Max Ports	
<div>Apply Reset</div>	

7. In *LA Port Settings*, verify that the member port is assigned to the port channel and the *Mode* is *Active*. *Port State* should be changed to *Up in Bundle* after the port link comes up.

<input type="checkbox"/>	Qx0/1	po100 ▾	Active ▾	128	Long ▾	2	Up in Bundle	Agg, Sync, Collect, Distrib.
--------------------------	-------	---------	----------	-----	--------	---	--------------	------------------------------

8. In *LA Interface Settings*, the port channel *Oper State* will become *Up* after at least one member port link comes up.

PORTCHANNEL INTERFACE BASIC SETTINGS

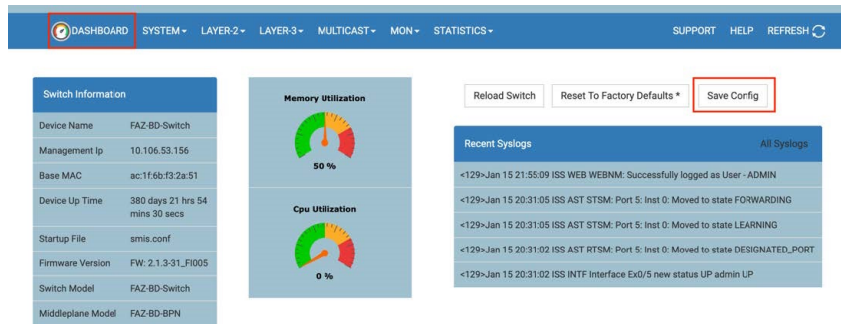
Port Channel ID	
Context ID	0 ▾
Admin Status	Up ▾
MTU	
<div>Add Reset</div>	

Select	Context ID	PortChannel ID	Admin State	Oper State	MTU
<input checked="" type="radio"/>	0	100	Up ▾	Up ▾	1500

Apply

Delete

9. Go to *Dashboard*, and then click *Save Config* to save the configuration.



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. See, [Initial set up on page 13](#).
2. Log in using the administrator credentials you created in [Set up Administrator accounts on page 15](#). 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 [Initial set up on page 13](#).



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 Security Event Manager cluster hosts using its internal IP. (For example, to SSH to the Controller host, use `exec ssh root@10.0.1.2`).

The IP is in can be determined by this format: `10.0.{chassis_id}.{blade_id}` where `10.0.*` 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 `fazbdctl` 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 `passwd` command.

GUI overview

FortiAnalyzer-BigData retains the same general GUI as the base FortiAnalyzer. In addition, there is a *Cluster Manager* tile.



Cluster Manager

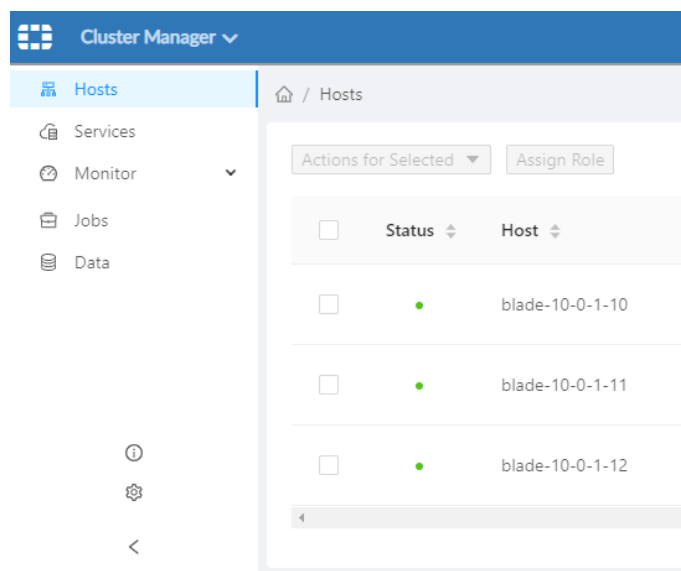
The Cluster Manager module enables you to manage hosts, services, logs, queries, jobs, and data resources in the Security Event Manager. See [Cluster Manager on page 32](#).

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 36 .
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 38 .
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 43 .
Jobs	The Jobs page manages system jobs and custom jobs. <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 80 .
Data	The Data page enables you to manages the data life cycle of your storage pools as well as data backups and restores. For more information, see Data management on page 88 .
System Information 	Click to see the current system version number.
System Upgrade 	Click to see your current system version and to upgrade FortiAnalyzer-BigData.

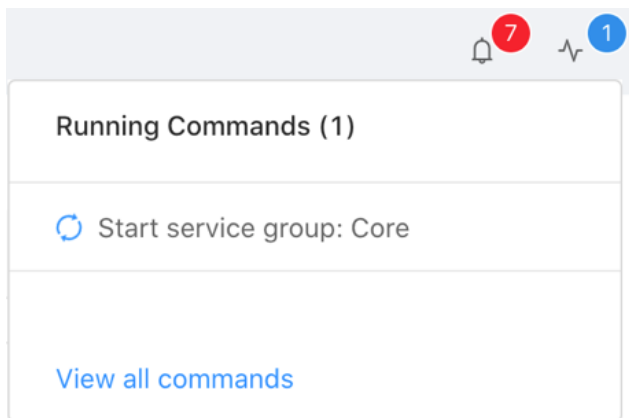
Custom refresh settings

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.

- Click *View all commands* at the bottom of the snapshot to view the full list of commands.

Running Commands

Start service group: Core

Recent Commands


Status	Command Name	Start Time		
✓	Start service group: Core	2/10/2020, 11:26:00 AM	📄	28.0s
✓	Start service group: Core	2/10/2020, 11:24:44 AM	📄	28.1s
✓	Assign role on worker-1	2/7/2020, 5:34:35 PM	📄	28.1s
✗	Restart instance: Query	1/30/2020, 1:36:52 PM	📄	4.3s
✗	Start service group: Data Lake	1/28/2020, 4:14:19 PM	📄	30.9s
✗	Start service group: Message Broker	1/28/2020, 4:14:18 PM	📄	29.8s
✗	Start service group: Core	1/28/2020, 4:14:16 PM	📄	31.4s

View all commands

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 Notification snapshot view and see more details. Clicking a notification item directs you to the page related to the notification event.

Unhealth Alerts

Unhealth Services 5

Unhealth Health Check Refresh 2

For the specific alerts such as the "Unhealth 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.

<input type="checkbox"/>	Status	Host Name	Role Type	Address	Instances	CPU Usage	Memory Usage	Disk Usage	Actions
<input type="checkbox"/>		blade-10-0-1-2	Master Node	10.0.1.2	▶ 9 Instance(s)	40.6%	40.7 GB / 125.6 GB	135.8 GB / 14.1 TB	Restart Status Details
<input type="checkbox"/>		blade-10-0-1-32	Master Node	10.0.1.32	▶ 7 Instance(s)	46.8%	50.5 GB / 125.6 GB	142.1 GB / 14.1 TB	Restart Status Details
<input type="checkbox"/>		blade-10-0-1-33	MetaStore Node	10.0.1.33	▶ 6 Instance(s)	35.6%	42.9 GB / 125.6 GB	132.3 GB / 14.1 TB	Restart Status Details
<input type="checkbox"/>		blade-10-0-1-34	MetaStore Node	10.0.1.34	▶ 9 Instance(s)	38.1%	59.1 GB / 125.6 GB	139 GB / 14.1 TB	Restart Status Details
<input type="checkbox"/>		blade-10-0-1-35	Data Node	10.0.1.35	▶ 5 Instance(s)	60.3%	52.5 GB / 125.6 GB	142.1 GB / 14.1 TB	Restart Status Details
<input type="checkbox"/>		blade-10-0-1-36	Data Node	10.0.1.36	▶ 5 Instance(s)	61.9%	45.7 GB / 125.6 GB	131.2 GB / 14.1 TB	Restart Status Details
<input type="checkbox"/>		blade-10-0-1-37	Data Node	10.0.1.37	▶ 5 Instance(s)	62.3%	52.7 GB / 125.6 GB	136.9 GB / 14.1 TB	Restart Status Details

The Host table contains the following columns:

Column header	Description
Status	There are three icons that represent the status of the host: <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 Name	The name of the host.
Role Type	Each host can have one of four roles. For more information about each role, see Roles on page 10 . <ul style="list-style-type: none">• Master Node• MetaStore Node• 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 37).
Address	The IP address of the host.
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.

Column header	Description
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"> • Restart: Restart the host. • Status Details: See the full metrics view of the host. • Assign Role: Assign a role to a new host.

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.

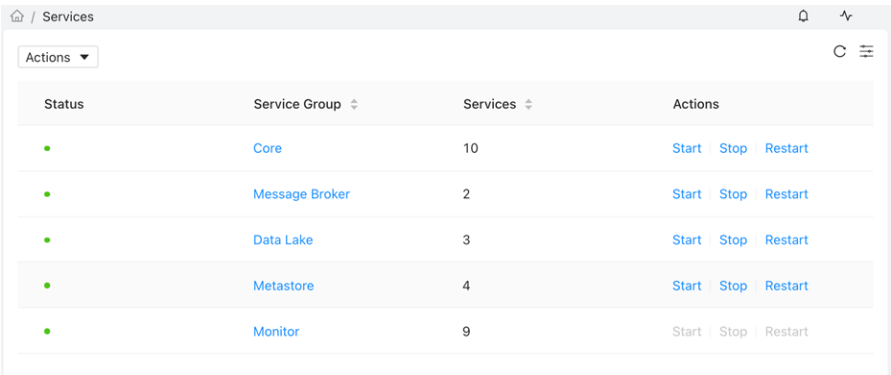


At this time, you can only assign the Data Node role.

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.



The screenshot shows the 'Services' page in the FortiAnalyzer interface. It features a table with four columns: Status, Service Group, Services, and Actions. The table lists five service groups: Core, Message Broker, Data Lake, Metastore, and Monitor. Each group has a green status icon, a count of services, and links to Start, Stop, and Restart actions.

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	<p>There are five icons that represent the status of the host:</p> <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	<p>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 39).</p>
Services	<p>The number of services running in each group.</p>
Actions	<p>There are three actions you can perform on each Service Group or service.</p> <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• Data Explorer• 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.

Home / Services / Data Lake

Instances Configuration

Actions for Selected

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.

Status	Service Name	Instances	Pending Configs	Actions
<input type="checkbox"/> ●	Kafka	8	0	Start Stop Restart

<input type="checkbox"/>	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 8

Data Directories:
log.dirs /data0/tmp/kafka-logs

* Data Retention Time:
log.retention.hours 6 Hours

* Data Retention Size:
log.retention.bytes 2000000001 Bytes ⓘ

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

* Number of Partitions:
num.partitions 26

* Default Replication Factors:
default.replication.factor 3

Kafka Heap Options: -Xmx8G -XX:G1HeapRegionSize=16M

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.

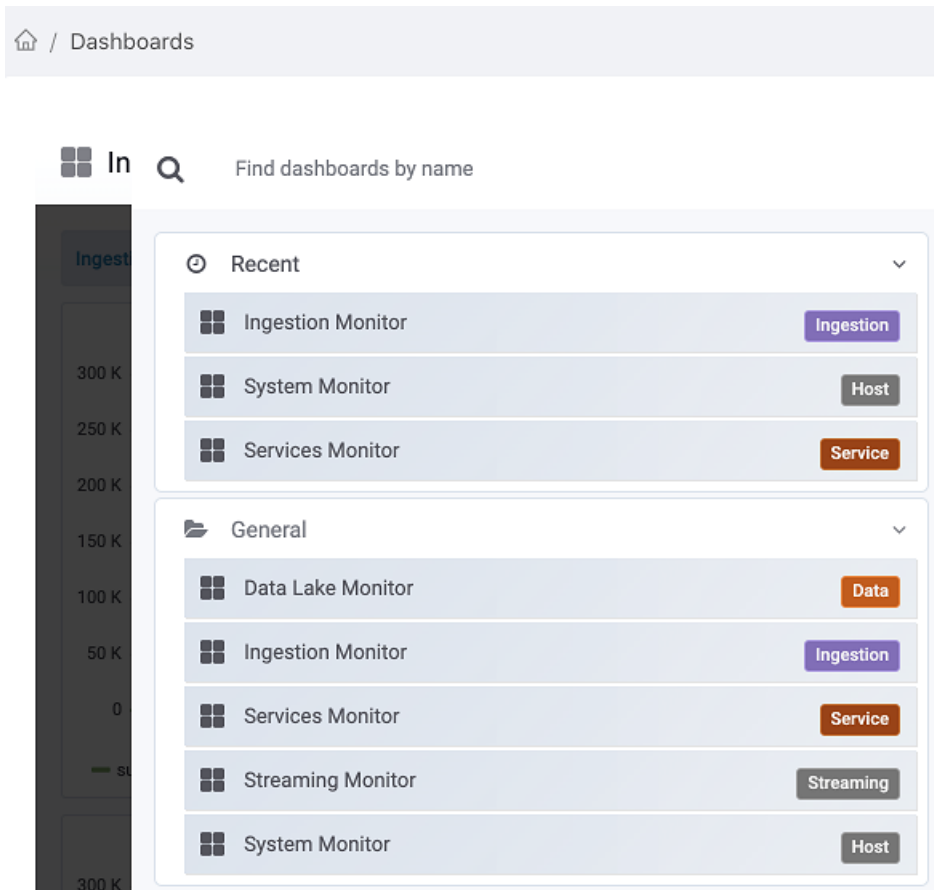
From the Dashboard, you can:

- 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.

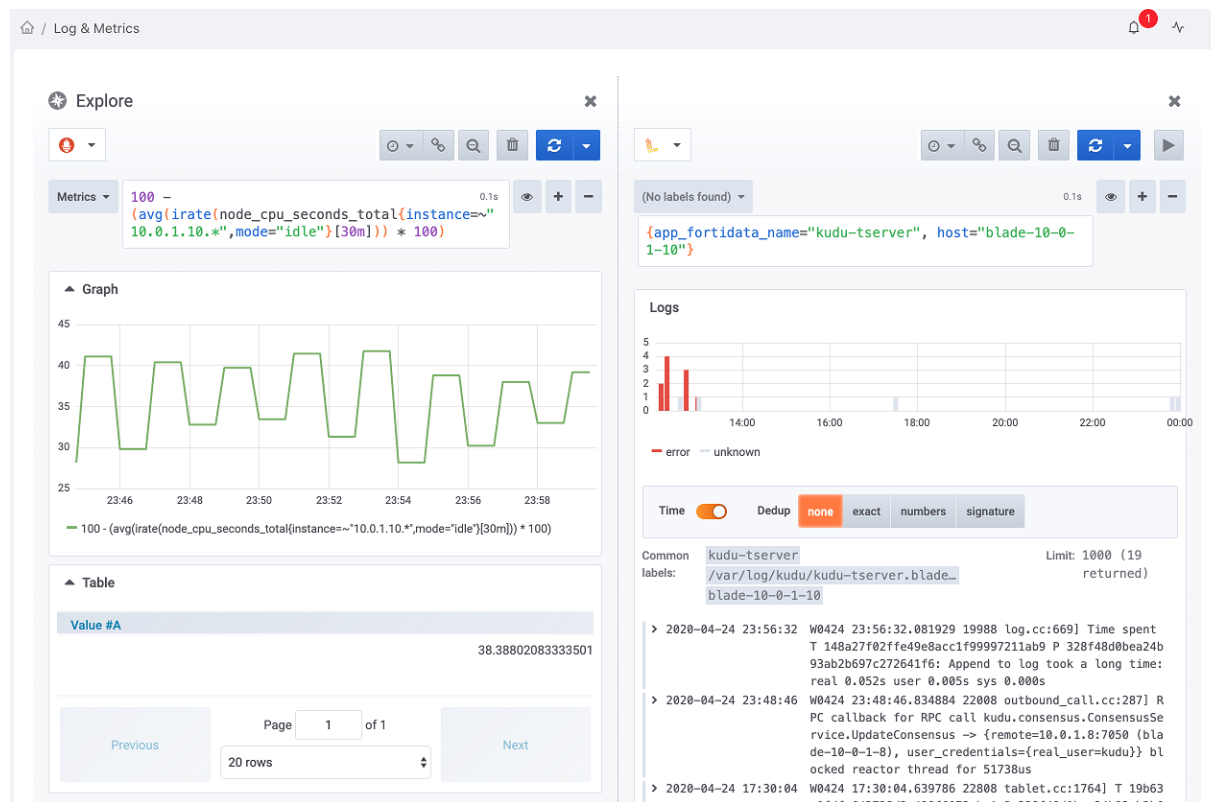
You can use the Explore search tool to switch between the Logs or Metrics view. The default selection is Logs.



- 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.



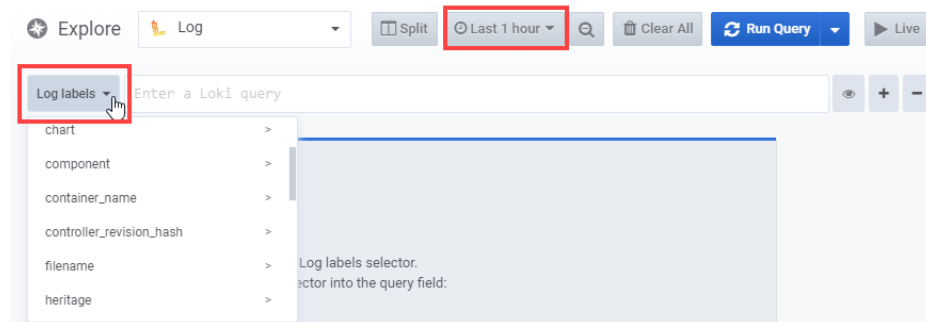
Explore logs

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"
```

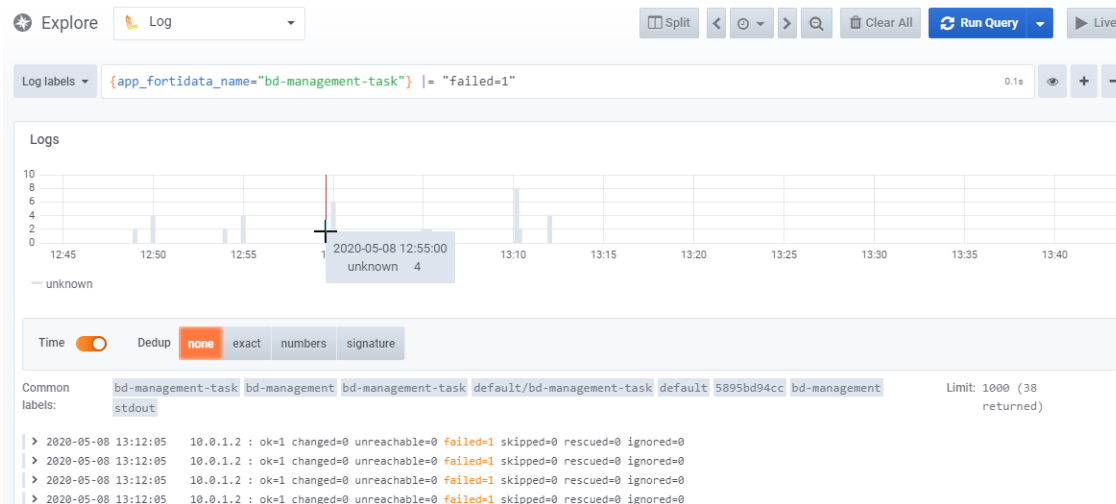
Supported operators:

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

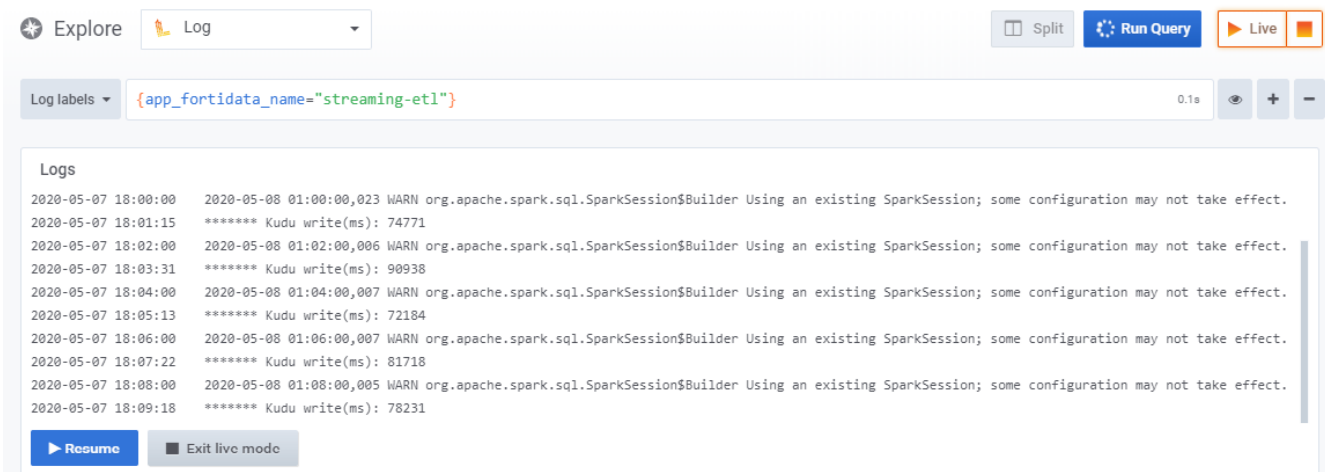
For more details, refer to the Loki query language (LogQL) documentation.

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.

The screenshot shows the FortiAnalyzer Monitor interface. At the top, there are tabs for 'Time', 'Dedup', 'none', 'exact', 'numbers', and 'signature'. Below these, there are common labels: 'bd-management-task', 'bd-management', 'bd-management-task', 'default/bd-management-task', 'default', '5895bd94cc', 'bd-management', 'stdout', and 'Limit: 1000 (38 returned)'. The main area displays a list of search results. Each result line starts with a timestamp and a status (e.g., 'ok=1 changed=0 unreachable=0 failed=1 skipped=0 rescued=0 ignored=0'). A red box highlights the 'Show context' link next to one of the results.

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

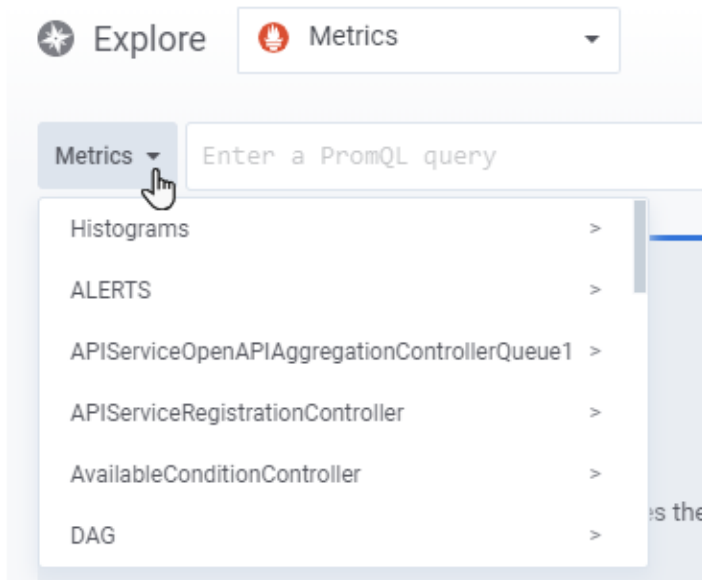
The screenshot shows the FortiAnalyzer Monitor interface with a context window open. The context window displays the following information:

- 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

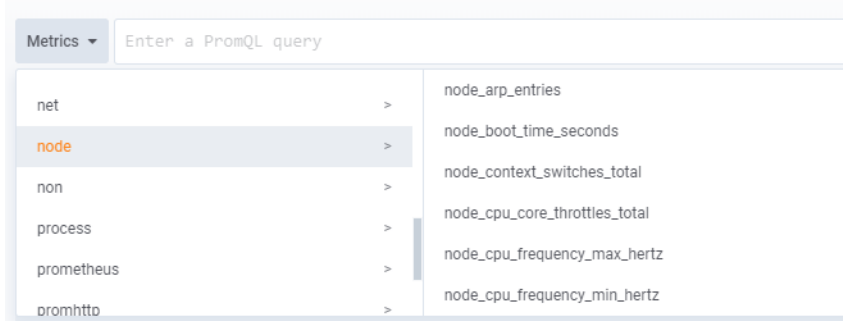
Explore metrics

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.

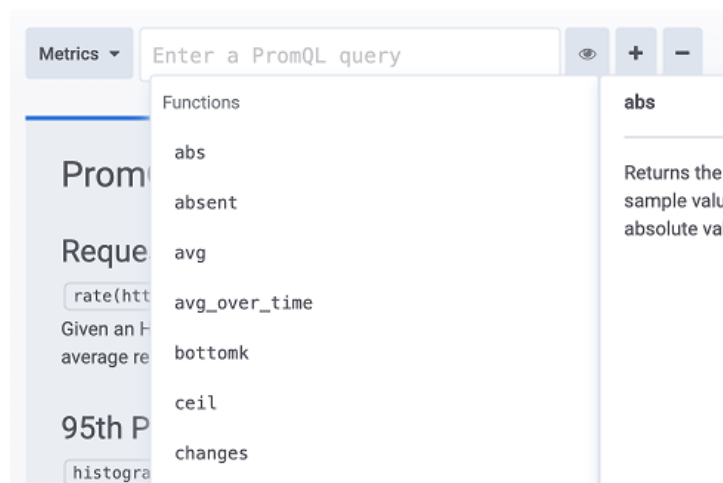


Metrics are grouped by prefixes, for example, all Node metrics are grouped under the "node" prefix.



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.

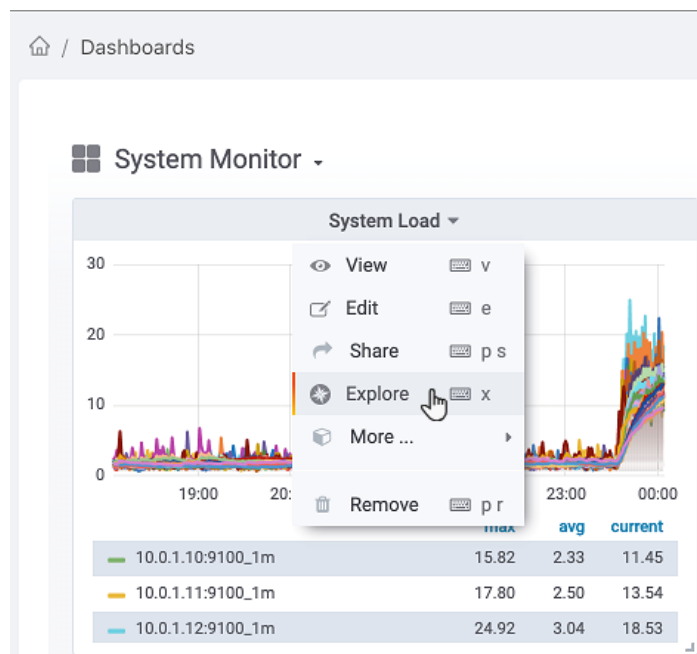


For more details, refer to the [Prometheus Query Language documentation](#).

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*.

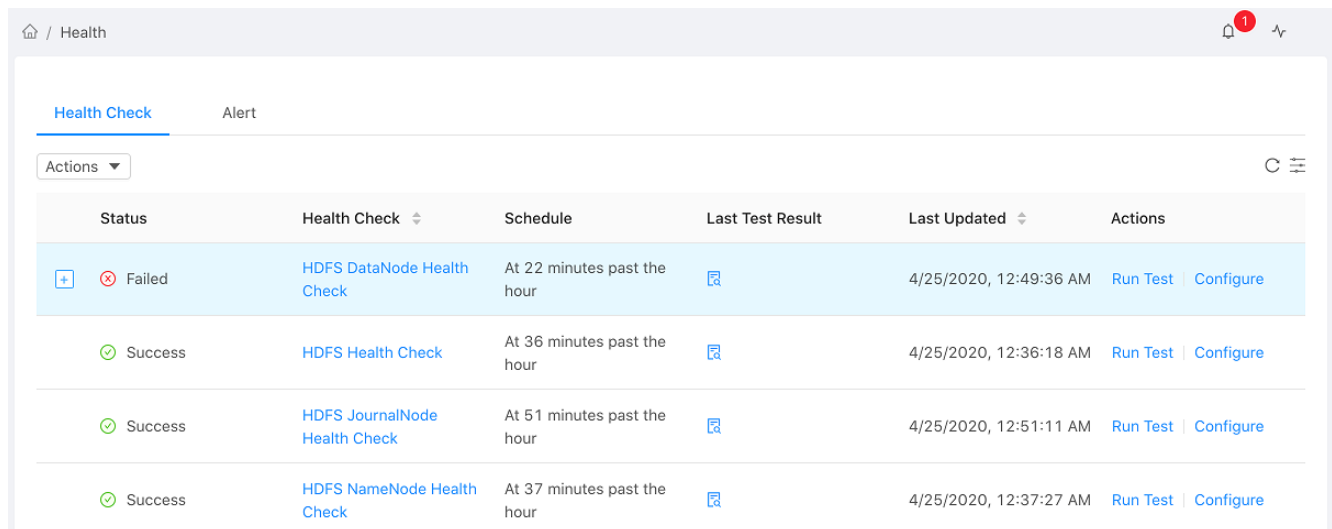


Health









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.



The screenshot shows the 'Health' page with the 'Health Check' tab selected. The table lists four predefined health checks. The first check, 'HDFS DataNode Health Check', is marked as 'Failed' with a red 'x' icon. The other three checks are marked as 'Success' with green checkmark icons. Each row includes a status icon, the health check name, the schedule, the last test result (with a document icon for details), the last updated timestamp, and two action links: 'Run Test' and 'Configure'.

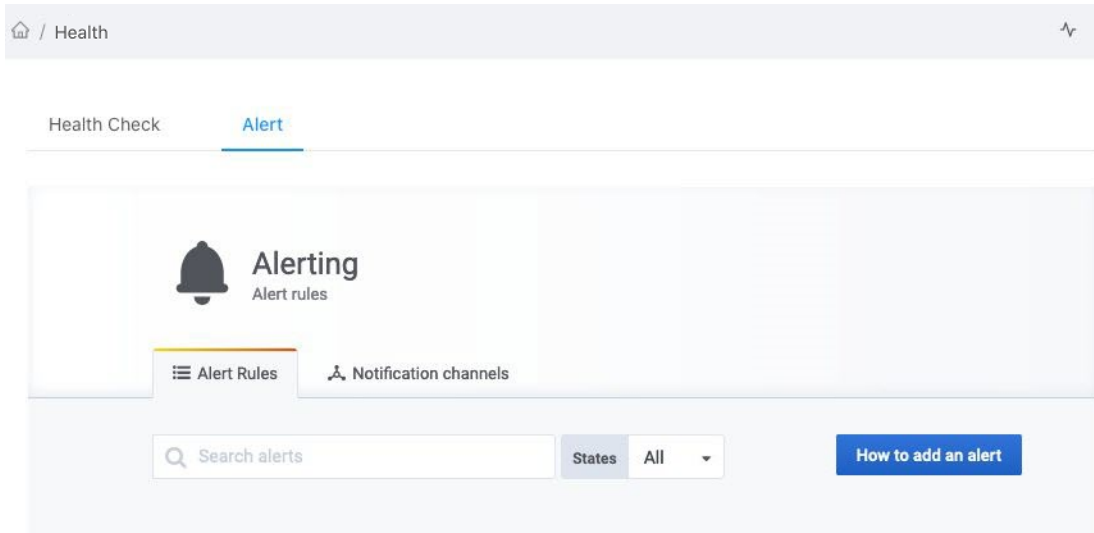
Status	Health Check	Schedule	Last Test Result	Last Updated	Actions
 Failed	HDFS DataNode Health Check	At 22 minutes past the hour		4/25/2020, 12:49:36 AM	Run Test Configure
 Success	HDFS Health Check	At 36 minutes past the hour		4/25/2020, 12:36:18 AM	Run Test Configure
 Success	HDFS JournalNode Health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

The Health Check table contains the following columns:

Column header	Description
Status	<p>Indicates if the health check was a success or failure.</p> <p>If a health check fails, you can click <i>Expand</i>  in the item row to see the error message.</p>
Health Test	<p>Shows what health check was run. You can click the name to see the history for that health check.</p> <hr/> <div>  <p>FortiAnalyzer-BigData only saves the last 500 records for each health check.</p> </div> <hr/>
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	<p>You can perform two actions on the health test:</p> <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 > Core > Configuration* and click *Monitor*.
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 Smtplib` command.
Wait for the command to finish running.
9. Go to *Monitor > Health > Alert > Notification channels* and click *Add 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*.

Home / Health

Health Check **Alert**

Edit Notification Channel

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

Slack settings

Url	https://hooks.slack.com/services/xxxxxxxxxxxxx...	
Recipient	#alerts	
Username		
Icon emoji		
Icon URL		
Mention		
Token	xoxb-0000000xxxxxxxxxxxxxxxx000000000000xox...	

Save
Send Test
Delete
Back

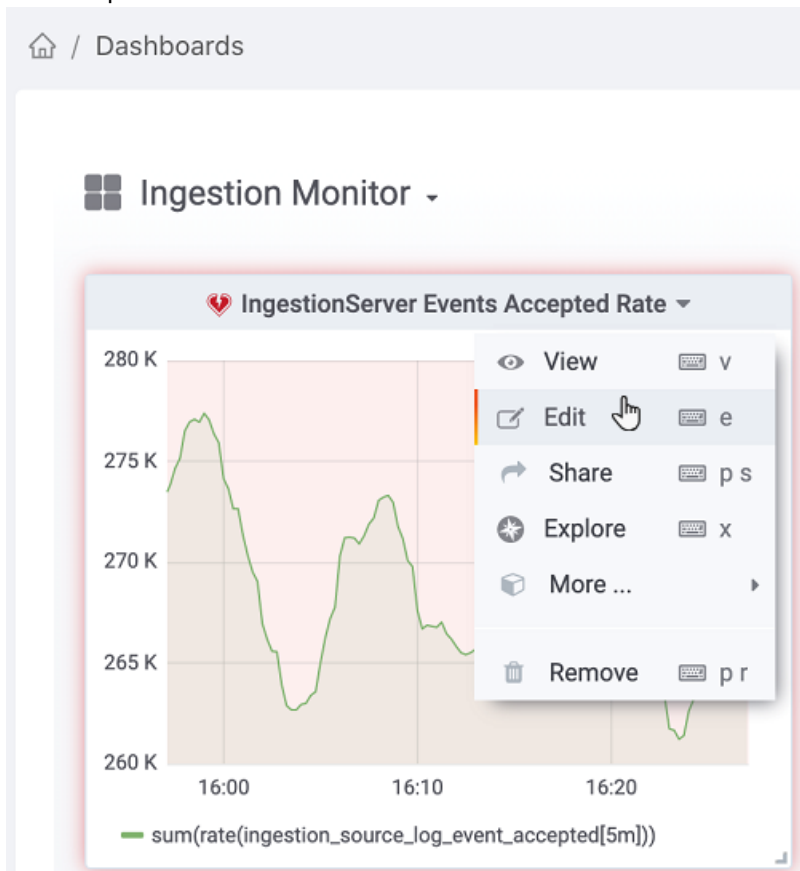
Custom alert rules

You can create custom alert rules from Dashboard panels and have it be 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 > Dashboard* and select a panel you want to create an alert for.
2. Click the panel title and click *Edit*.



The panel's detailed view loads.

3. Click *Alert*  to access the Alert view and click *Create Alert* to specify conditions that trigger the alert.

4. You can create conditions through two different methods:

- By making queries in the Conditions section.

The screenshot shows the 'Alert' configuration page. At the top, there are buttons for 'State history', 'Test Rule', and 'Delete'. The 'Name' field is 'IngestionServer Events Accepted Rate a...'. The 'Evaluate every' field is '1m' and the 'For' field is '1m'. Under the 'Conditions' section, there is a 'WHEN' clause with 'avg ()' and an 'OF' clause with 'query (A, 5m, now) IS ABOVE 250000'. There is a '+' button to add more conditions. Under the 'No Data & Error Handling' section, there are two rows: 'If no data or all values are null' with 'SET STATE TO' 'No Data', and 'If execution error or timeout' with 'SET STATE TO' 'Alerting'.

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



5. After you've defined your condition, select the Notification Channel and click *Test Rule* to test the alert rule.

6. Click **Save** to save your settings.

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


FortiOPS APP


[Alerting] IngestionServer Events Accepted Rate alert

Ingestion Rate is above threshold.

`{}`

268109.18819735

 Grafana v6.5.2 | Today at 3:54 PM



Hyperscale firewall NetFlow logging support

You can configure FortiAnalyzer-BigData as a NetFlow 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, 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 IPFIX log messages, 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 IPFIX log message traffic.

To set up external IP addresses for Security Manager hosts:

1. Go to *Cluster Manager > Hosts* and click *External IPs*.
2. Set the *Default Gateway* and *Netmask* to your internal network.

3. Enter an IP address for each blade.

The screenshot shows a web interface for configuring NetFlow logging. At the top, there's a breadcrumb 'Home / Hosts' and a toolbar with 'Actions for Selected', 'Assign Role', and 'External IPs'. The main dialog is titled 'Set the external IPs for NetFlow log traffic'. It has two tabs: 'Set by Host' (selected) and 'Set with range'. Under 'Set by Host', there are input fields for 'Default Gateway' (10.106.2.254), 'Netmask' (255.255.255.0), and a list of blades with their respective IP addresses: blade-10-0-1-10 (10.106.2.137), blade-10-0-1-11 (10.106.2.138), blade-10-0-1-12 (10.106.2.139), blade-10-0-1-13 (10.106.2.140), blade-10-0-1-14 (10.106.2.141), blade-10-0-1-2 (10.106.2.129), blade-10-0-1-3 (10.106.2.130), and blade-10-0-1-4 (10.106.2.131). At the bottom right 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.

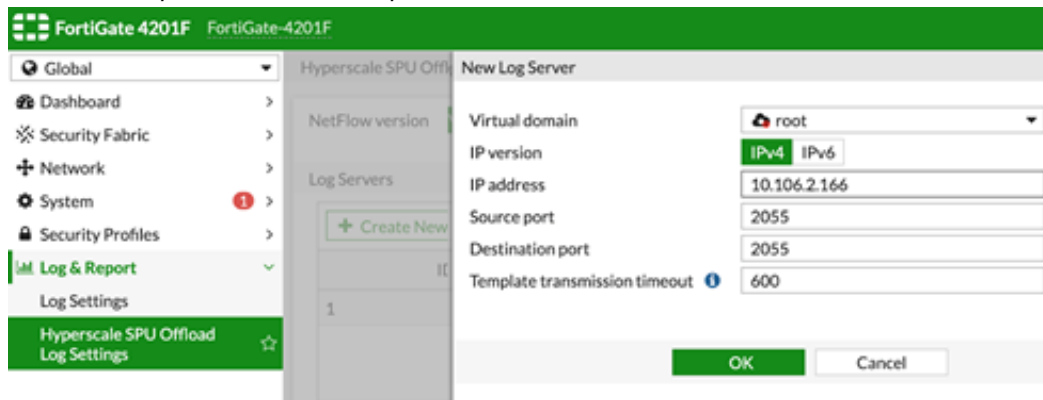
The screenshot shows the same web interface as before, but the 'Set with range' tab is selected. The 'Start External IP' field is set to 10.106.2.139. The 'Netmask' field is 255.255.255.0 and the 'Gateway' field is 10.106.2.254. The 'Cancel' and 'Apply' buttons are at the bottom right.

Configure FortiAnalyzer-BigData as IPFIX log server on 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) log messages over UDP to FortiAnalyzer-BigData. For more information, see [Hyperscale Firewall Hardware logging](#) in the [Fortinet Doc Library](#).

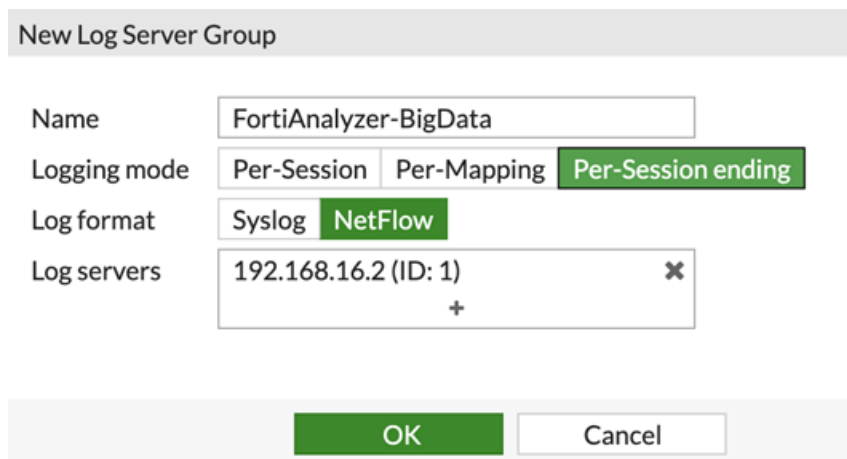
To configure FortiAnalyzer-BigData as NetFlow log server on FortiGate:

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.



Device Manager and log rate

When FortiAnalyzer-BigData starts to receive IPFIX 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 IPFIX logs, go to *Cluster Manager > Monitor > Dashboards* to view the *Ingestion Log Accepted Rate*.



Search IPFIX log in Log View

If you have IPFIX logs collected, you can go to *FortiGate > IPFIX* 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 left sidebar contains navigation options: FortiGate, Traffic, Security (Antivirus, Intrusion Prevention, Application Control, Web Filter, DNS, Vulnerability Scan, VoIP), Event, GTP, IPFIX (selected), FortiClient, FortiAnalyzer, Custom View, Log Browse, and Log Group. The main area displays a table of log entries for IPFIX logs. The table has columns: #, Date/Time, Device ID, Virtual Domain, Source Address, and Post NAT Source Address. The table shows 11 log entries. Below the table, there is a 'logDetails' section showing expanded information for a selected log entry (entry 1).

#	Date/Time	Device ID	Virtual Domain	Source Address	Post NAT Source Address
1	22:51:32	FG441FTK20900185	root	1.1.1.1	1.1.1.102
2	22:51:32	FG441FTK20900185	root	2018:1::2	2018:1::2
3	22:51:32	FG441FTK20900185	root	2018:1::2	2018:1::2
4	22:51:32	FG441FTK20900185	root	1.1.1.1	1.1.1.102
5	22:51:32	FG441FTK20900185	root	2018:1::2	2018:1::2
6	22:51:32	FG441FTK20900185	root	2018:1::2	2018:1::2
7	22:51:32	FG441FTK20900185	root	2018:1::2	1.1.1.101
8	22:51:32	FG441FTK20900185	root	1.1.1.5	64:ff9b::101:105
9	22:51:32	FG441FTK20900185	root	1.1.1.5	64:ff9b::101:105
10	22:51:32	FG441FTK20900185	root	2018:1::2	1.1.1.101
11	22:51:32	FG441FTK20900185	root	2018:1::2	1.1.1.101

Below the table, the 'logDetails' section shows expanded information for a selected log entry (entry 1):

- CVE ID
- Date/Time: 22:51:32
- Device ID: FG441FTK20900185
- Device Name: FG441FTK20900185
- NAT Event: 17
- Port Range End: 128
- Port Range Start: 1
- Post NAT Source Transport...
- Post NAT Source Address: 1.1.1.101
- Protocol Identifier: 17
- Source Address: 2018:1::2
- Source Transport Port: 2233
- Template ID: 285
- Time Stamp: 2021-05-14 22:51:32
- Virtual Domain: root
- exportingProcessId: 2323644417

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 64](#)
- [Log types \(Global Search\) on page 65](#)
- [Create a new Search Federation \(Example\) on page 71](#)
- [Log Query Language \(LogQL\) on page 77](#)

Starting a global search

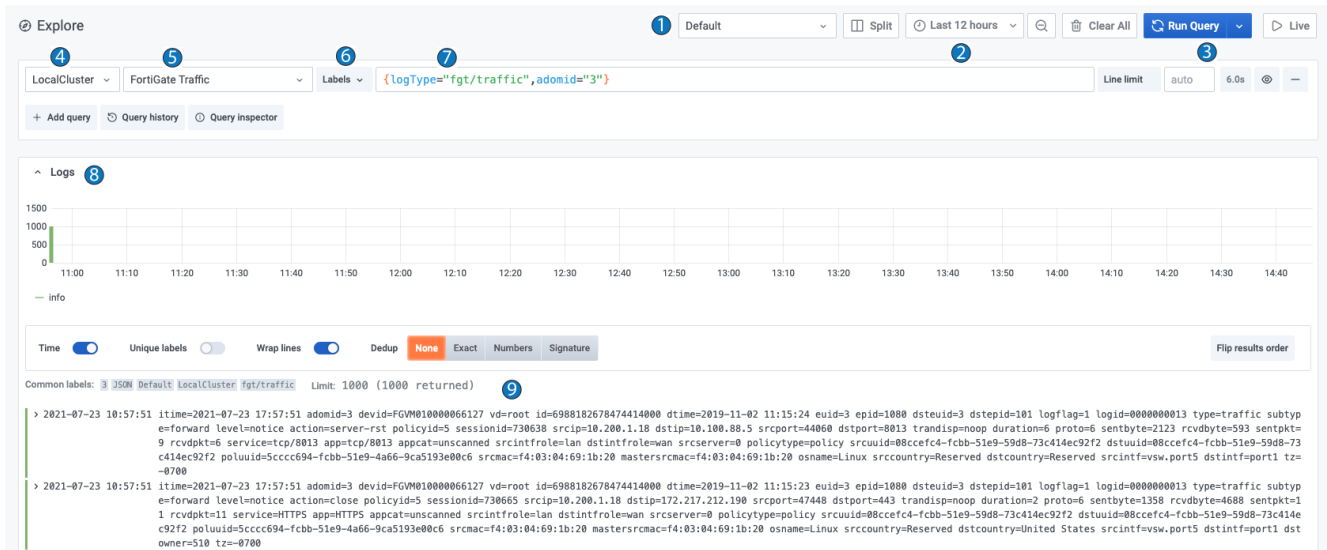
To perform a simple global search, use the default values for the search federation, cluster and log time, then select a log type.

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.

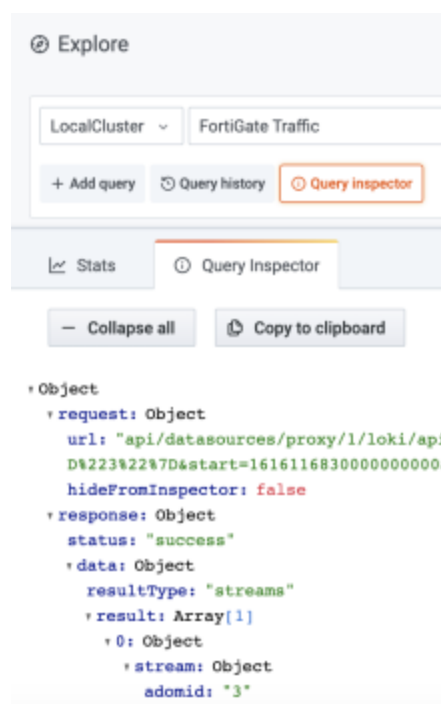
Histogram

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.



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:

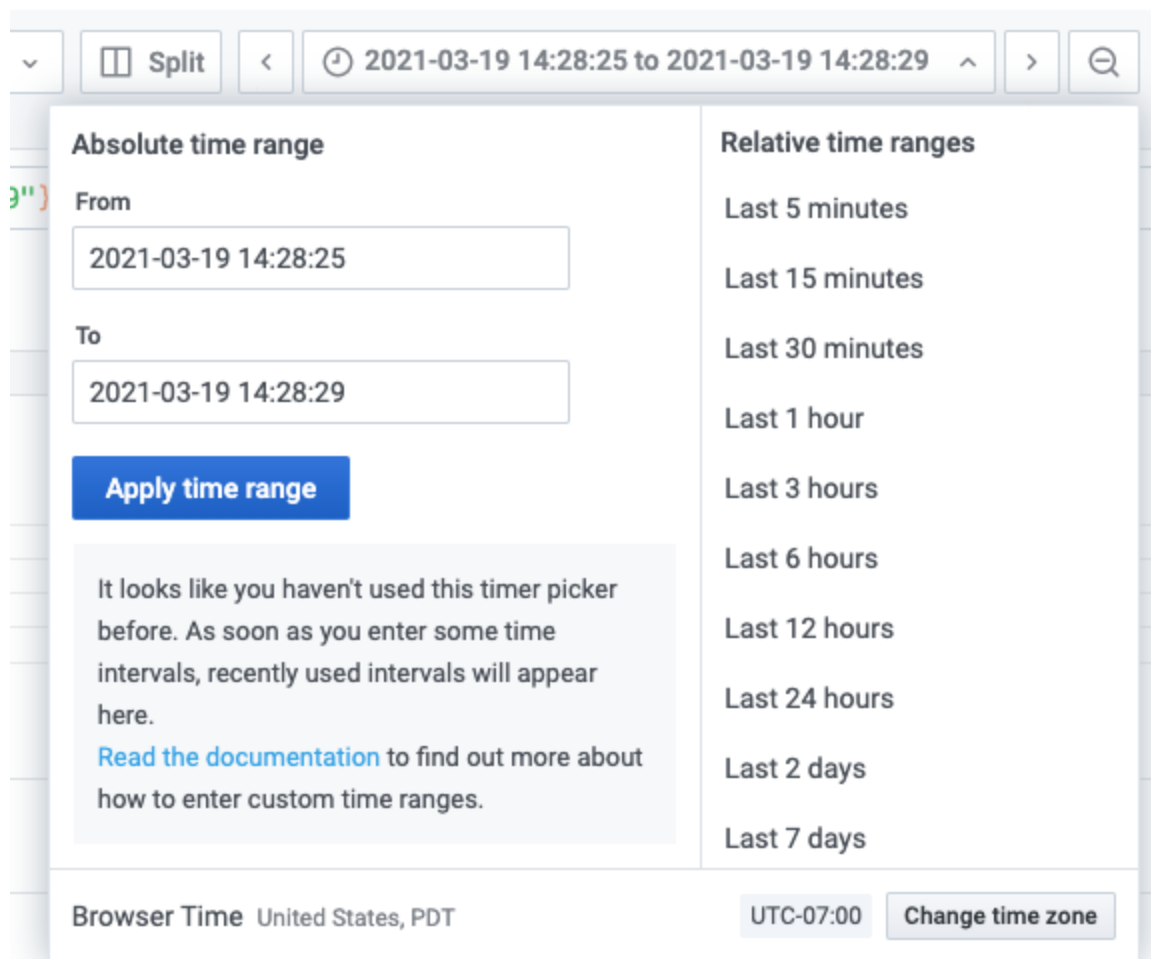
Icon	Description
	Shows the value statistics of the label or fields.
	Adds the label or field as a condition in query language. For example, app=DNS.
	Excludes the label or field in the query language. Fore example, app!=DNS.
	Displays only this label or field in the log item. For example, app=DNS.

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 the Time Window interface. At the top, there is a search bar with a dropdown arrow, a 'Split' button, a left arrow, a clock icon, a time range '2021-03-19 14:28:25 to 2021-03-19 14:28:29', an up arrow, a right arrow, and a magnifying glass icon. Below this, there are two columns: 'Absolute time range' and 'Relative time ranges'. The 'Absolute time range' column has 'From' and 'To' fields with the same time range as the search bar, and an 'Apply time range' button. The 'Relative time ranges' column lists various 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, there is a 'Browser Time' section showing 'United States, PDT' and 'UTC-07:00', with a 'Change time zone' button.

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.

Explore Default 🔍 🗑️ 🔄 ▶️

LocalCluster FortiGate Traffic Labels `{adomid="3",appcat="Web.Client",logType="fgt/traffic"}` Line limit auto 3.9s 🔍 —

[+ Add query](#) [Query history](#) [Query inspector](#)

[Query history](#) [Starred](#) [Settings](#) ×

Filter history Newest first

today

March 18 3 queries

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

a week ago

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.

Explore Default 🔍 🗑️ 🔄 ▶️

LocalCluster FortiGate Traffic Labels `{logType="fgt/traffic",adomid="3"}` Line limit auto 3.4s 🔍 —

[+ Add query](#) [Query history](#) [Query inspector](#)

Logs

Time ☒ Unique labels ☐ Wrap lines ☒ Dedup **None** Exact Numbers Signature

Flip results order

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=FGVM01000066103 vd=root id=694148548763477
6000 dtime=2019-11-02 02:30:19 euid=3 epid=1065 dstuid=3 dstpid=101 logflag=1 logi
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=46938 dstport=53 transp=snmp
duration=181 proto=17 sentbytes=116 rcvbytes=16 sentpkt=2 rcvpkt=2 service=DNS app
DNS appcat=unscanned srcintfrole=lan dstintfrole=wan srcserver=0 polycity=policy 5
rcuid=08ccefca-fcbb-51e9-59d8-73c414ec92f2 dstuid=08ccefca-fcbb-51e9-59d8-73c414ec
92f2 poluid=5ccc694-fcbb-51e9-4a66-9ca5193e00c6 srcmac=00:14:c2:26:92:82 mastersrc
```

Explore Default 🔍 🗑️ 🔄 ▶️

LocalCluster FortiGate DNS Labels `{logType="fgt/dns",adomid="3",serverName="LocalCluster"}` Line limit auto 2.2s 🔍 —

[+ Add query](#) [Query history](#) [Query inspector](#)

Logs

Time ☒ Unique labels ☐ Wrap lines ☒ Dedup **None** Exact Numbers Signature

Flip results order

Common labels: 3 JSON Default LocalCluster fgt/dns Limit: 835 (835 returned)

```
> 2021-03-19 14:47:49 itime=2021-03-19 21:47:49 adomid=3 devid=FGVM01000066102 vd=root id=694148520846190
2000 dtime=2019-11-02 02:29:57 euid=3 epid=1044 dstuid=3 dstpid=101 logid=15010548
03 type=utm 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=50460 at
ypeval=1 srcintfrole=lan dstintfrole=wan ipaddr=(200.91.112.55) srcintfport3 dstint
fport1 profile=default qname=99.goodyouxi.com qtype=A qclass=IN catdesc=Malicious W
eb sites 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.

The screenshot shows the FortiAnalyzer interface for Live Log search. At the top, there's a search bar with a query: `{logType="fgt/traffic", adomid="3"}`. To the right of the search bar are buttons for 'Run Query', 'Live', and 'Pause'. Below the search bar, there are tabs for 'LocalCluster', 'FortiGate Traffic', and 'Labels'. The 'Labels' tab is selected, showing a query: `{logType="fgt/traffic", adomid="3"}`. The search results are displayed in a log format, showing two entries for 2021-03-19 15:00:16. The first entry is a traffic log with details like srcip, dstip, srcport, and dstport. The second entry is a threat log with details like threattype, threatname, and threatid. The interface also includes a 'Pause' button and an 'Exit live mode' button.

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

Incoming Federation Request

☐ Allow incoming federation request

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

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


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

X

* Target Server Name:

10.105.101.59

* Target Server Address:

10.105.101.59

Target User Name:

Auto-confirm Token:

NOTE: This feature only works when sites
certificates are trusted

Cancel

OK

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
-------------	--

Deny
Accept

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

3. 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						
<input checked="" type="checkbox"/> 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


4. 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

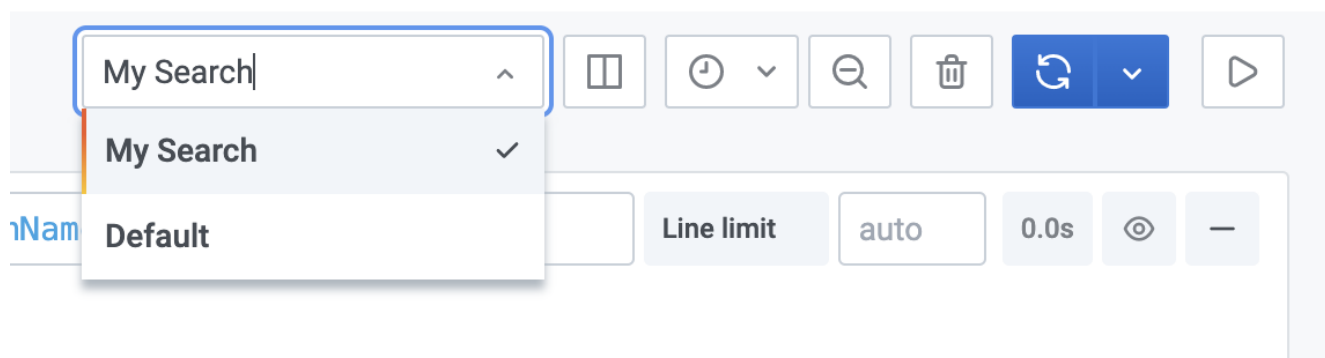
* Federation Name:

Federation Servers: ☒ 10.105.101.4
☒ 10.105.101.59

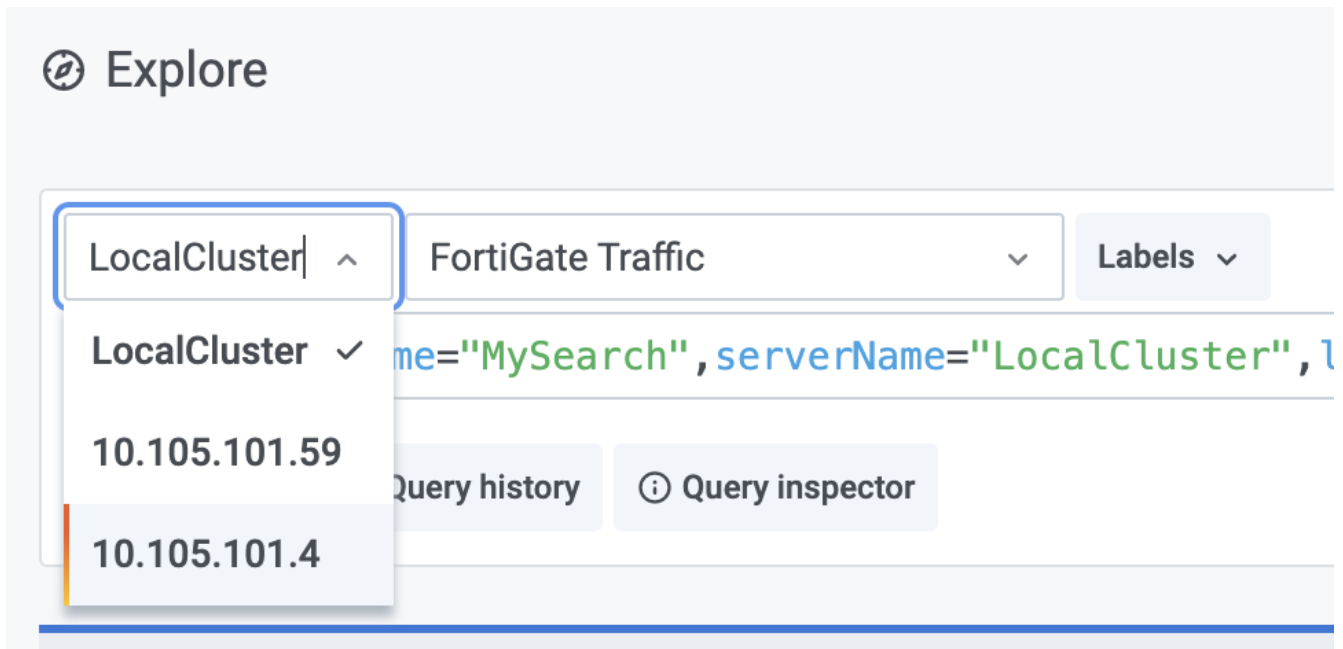
Search Federation			
+ Add Federation			
Federation Name	Created At	Last Updated	Actions
 My Search	2/2/2021, 8:14:33 PM	2/2/2021, 8:14:33 PM	Edit Delete
Server Name		Server Address	
• 10.105.101.4		10.105.101.4	
• 10.105.101.59		10.105.101.59	

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, • <code>*</code> represents a sequence of zero or more items, • <code>+</code> represents a sequence of one or more items, • <code>?</code> produces a non-greedy match, and so on.

Operator	Description
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 buttons for '+ 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 81).
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	<p>You can perform two actions on the health test:</p> <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.

Home / 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 500 records for job execution results

Built-in automation jobs

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

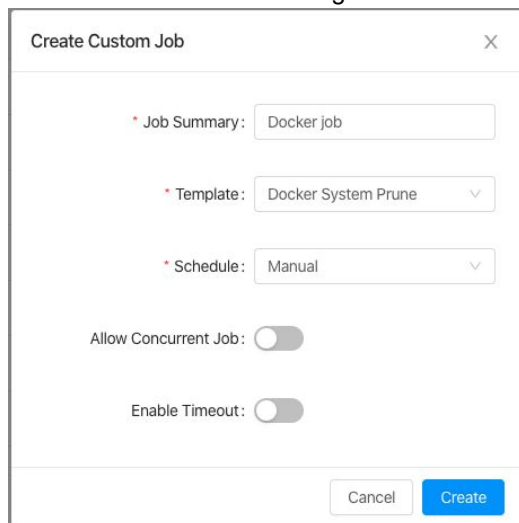
Built-in jobs	Description
Data Retention	Created automatically when storage pools are created. This job is used to apply data retention policies for the Storage Pool which marks the old data for deletion and makes space for future data.
Data Rebalance	Created automatically when storage pools are created. This job is used to rebalance Kudu data partitions to evenly distribute them across the Security Event Manager hosts.
Data Appendix	Created automatically when storage pools are created. This job generates the list of available sub-types of FortiGate Event logs for LogView.
Facet Formation - Report	Created automatically when storage pools are created. This job generates the pre-aggregated facets to speed up FortiView queries.
Facet Formation - FortiView	Created automatically when storage pools are created. This job generates the pre-aggregated facets to speed up Report queries.
Storage Pools Restore	This job will be created automatically when you launch the storage pool restore function from the Data page. For more details, see Data restore on page 96 .

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.



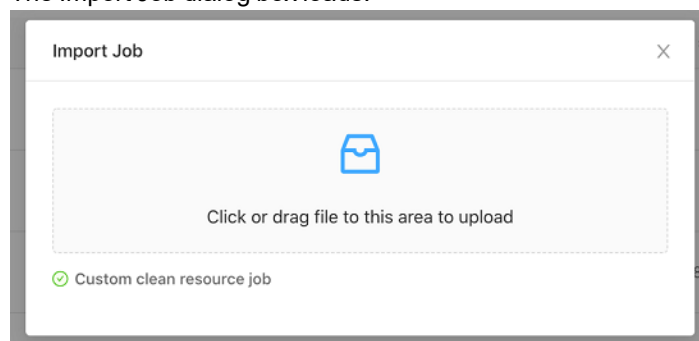
2. Complete the following fields:

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 84 .
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 | These inventory groups can be used to select the host(s) that have the named services running. |
|--|--|

<ul style="list-style-type: none"> • kudu_hive_metastore • zookeeper • kafka_broker • impala_catalog • impala • impala_statestore • yarn_nodemanager • yarn_resource_manager • spark_history_server 	For example, using “ host: kudu_tserver” in your playbook allows it to be executed on all hosts has kudu-tserver instance.
<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 • yarn_resource_manager_reachable • spark_history_server_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"> • 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 116](#)).

Kafka Rebalance

The Kafka Rebalance template rebalances the data load across the Security Event Manager 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 108](#).

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 116](#)).

Purge Data Pipeline

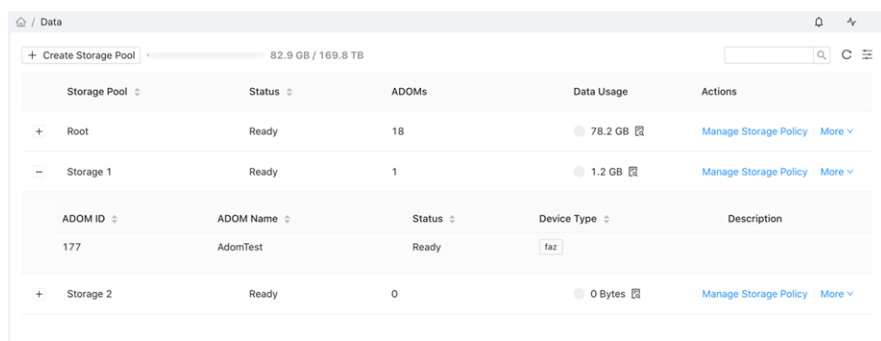
This job resets the watermark and performs a clean restart of the pipeline.



Any unprocessed data will be lost (see [How to recover from an unhealthy service status on page 116](#)).

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	Status	ADOMs	Data Usage	Actions
Root	Ready	18	78.2 GB	Manage Storage Policy More
Storage 1	Ready	1	1.2 GB	Manage Storage Policy More

ADOM ID	ADOM Name	Status	Device Type	Description
177	AdomTest	Ready	faz	

Storage 2	Ready	0	0 Bytes	Manage Storage Policy More
-----------	-------	---	---------	--

The Storage Pool table contains the following columns:

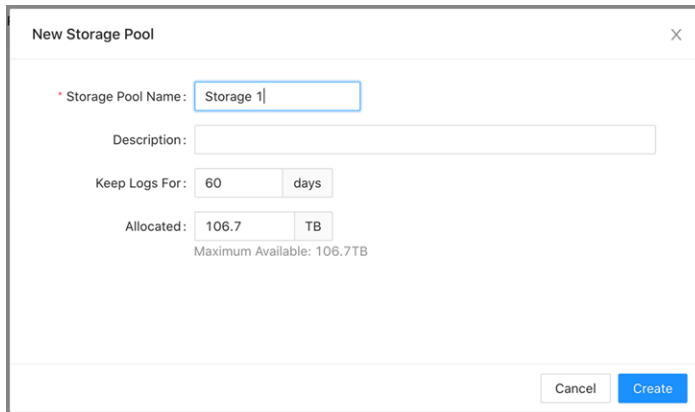
Column header	Description
Storage Pool	The name of storage pool. You can expand each storage pool to display all the ADOMs in that group.
Status	Indicates the status of the storage pool. <ul style="list-style-type: none">• Ready: The storage pool is ready for use.• In Progress: The storage pool is being created and is not yet ready for use.• Failed: The storage pool creation failed.
ADOMs	The number of ADOMs in that storage pool.
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 90.• Manage Job: Manage jobs in that storage pool.• Backup: Create a backup of that storage pool.• Restore: Restore data.

To create a storage pool:

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

Storage Pool Name	Enter a name for the storage pool.
--------------------------	------------------------------------

Description	Enter a description of the storage pool.
Keep Logs For	Enter the number days the logs are to be stored.
Allocated	Enter the amount of memory allocated for storing the logs.



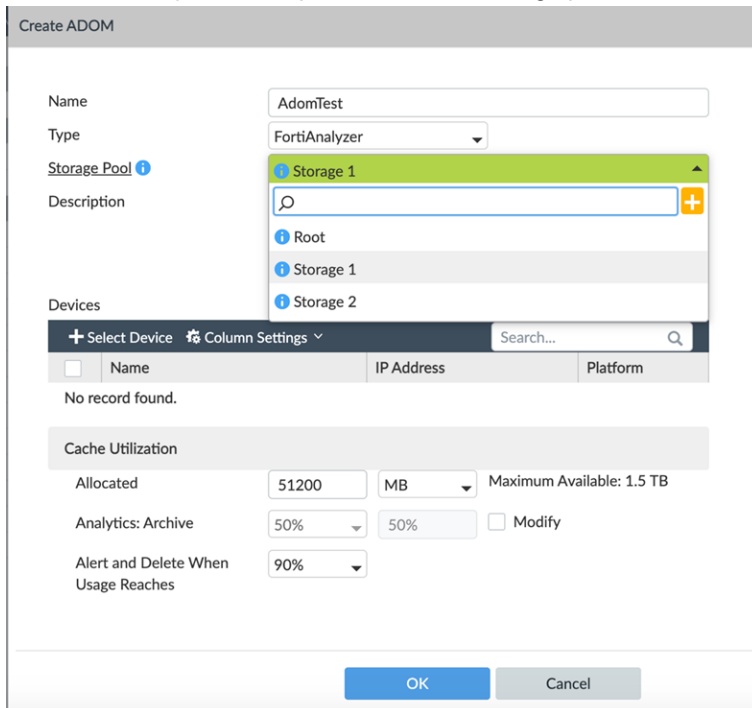
The 'New Storage Pool' dialog box contains the following fields and controls:

- Storage Pool Name:** A text input field with 'Storage 1' entered.
- Description:** An empty text input field.
- Keep Logs For:** A numeric input field set to '60' and a unit dropdown menu set to 'days'.
- Allocated:** A numeric input field set to '106.7' and a unit dropdown menu set to 'TB'.
- Maximum Available:** A label indicating '106.7TB'.
- Buttons:** 'Cancel' and 'Create' buttons at the bottom right.

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.



The 'Create ADOM' dialog box contains the following fields and controls:

- Name:** A text input field with 'AdomTest' entered.
- Type:** A dropdown menu set to 'FortiAnalyzer'.
- Storage Pool:** A dropdown menu showing 'Storage 1' as the selected option. Below it is a search bar and a list of options: 'Root', 'Storage 1', and 'Storage 2'. A plus sign (+) is visible next to the search bar.
- Description:** An empty text input field.
- Devices:** A section with a '+ Select Device' button, a 'Column Settings' dropdown, and a search bar. Below this is a table with columns: Name, IP Address, and Platform. The table is currently empty, with the text 'No record found.' below it.
- Cache Utilization:** A section with the following controls:
 - Allocated:** A numeric input field set to '51200' and a unit dropdown menu set to 'MB'. A label 'Maximum Available: 1.5 TB' is next to it.
 - Analytics: Archive:** A dropdown menu set to '50%' and a 'Modify' checkbox.
 - Alert and Delete When Usage Reaches:** A dropdown menu set to '90%'.
- Buttons:** 'OK' and 'Cancel' buttons at the bottom.

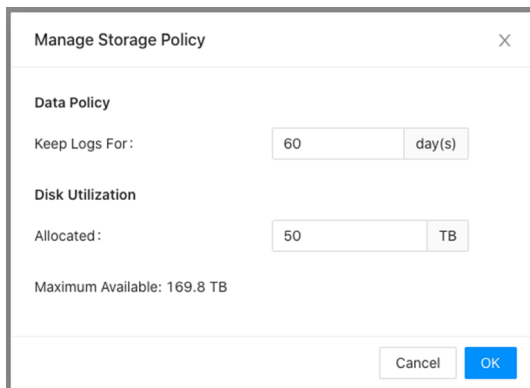
Manage storage policy

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

To manage a storage policy:

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

The *Manage Storage Policy* dialog opens:



The dialog box titled "Manage Storage Policy" contains the following fields:

- Data Policy**
 - Keep Logs For: 60 day(s)
- Disk Utilization**
 - Allocated: 50 TB
 - Maximum Available: 169.8 TB

At the bottom right are "Cancel" and "OK" buttons.

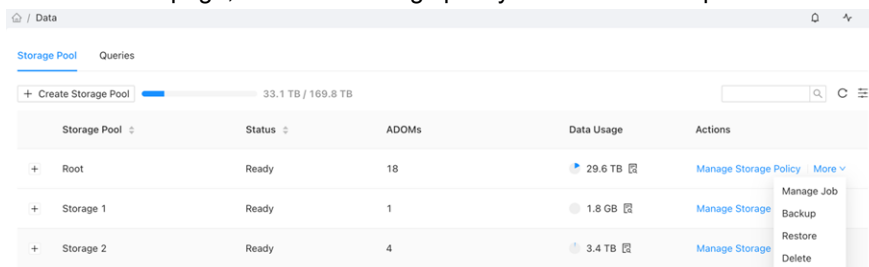
2. In the *Keep Logs For* field, select the number of days you want to store the data in the system. FortiAnalyzer-BigData removes the data from the system after the selected number of days.
3. In the *Allocated* field, select the maximum amount of data you want to store in the system. FortiAnalyzer-BigData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

Data backup

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 screenshot shows the "Data" page with a table of storage pools. The "Actions" column for "Storage 2" is expanded, showing the "Backup" option.

Storage Pool	Status	ADOMs	Data Usage	Actions
Root	Ready	18	29.6 TB	Manage Storage Policy, More
Storage 1	Ready	1	1.8 GB	Manage Storage
Storage 2	Ready	4	3.4 TB	Manage Storage, Backup, Restore, Delete

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.</p> <hr/> <div>  <p>If the URL is configured to an external HDFS cluster, all its hosts must be made accessible by the Security Event Manager hosts (see Backup and restore to external HDFS on page 103).</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>
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>  <p>This mode may take longer to complete the back up, so only enable Safe Mode when the normal backup job fails.</p> </div>
Advanced Config	<p>These configurations define the resources used for the job. Normal users should keep the default configurations.</p>
Enable Scheduled Backup	<p>Enable so the backup can be scheduled automatically.</p>

- 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*.

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

✕

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: ☐

Cancel

Save

Save & 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.

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

✕

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: ☐

Cancel
Save
Save & 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:

Storage backup at Tue Apr 28 23:00:00 GMT 2021

▼

* HDFS URL:

hdfs://cluster/backup001

Backup Tables:

Backup Timestamp:

11-18-2021 19:09:42

Restore Timeout:

Hour(s)

Enable Safe Mode:

?

Advanced Config:

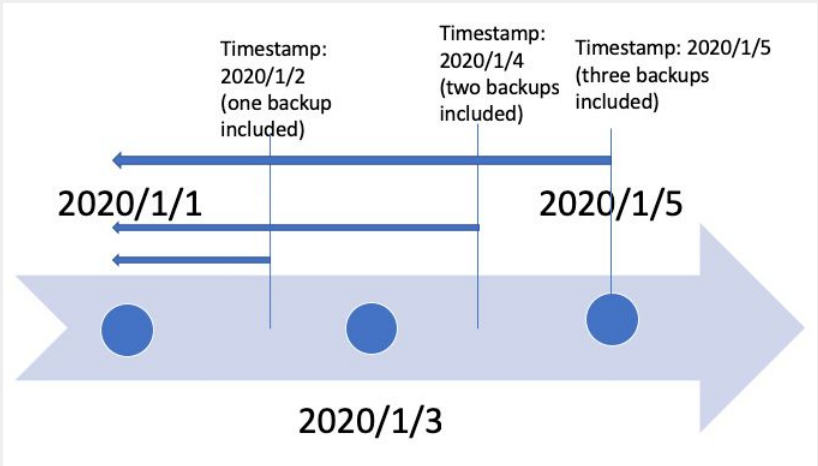

Restore storage pool action will pause the log input (Ingestion and Pipeline will be stopped), and **DELETE** all data (**CANNOT** be undone) in current storage pool. Please type **restore** to confirm:

Cancel

Restore

- 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.

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 BigData hosts (Blade A2-A13) except the Main host (Blade A1).



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

1. Connect to the CMM web management utility (see [Connect to the Chassis Management Module on page 16](#)).
2. Select one of the Security Event Manager host (see [Remotely control blades via CMM on page 22](#)) to enter its bootloader.
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. Immediately after you reboot a host, press the **Tab** key within 10 seconds to bring out the action options.
7. When the following options show up, type `bootloader` to enter the bootloader's main page.

```
SYSLINUX 4.05 EDD 0x5bd8f633 Copyright (C) 1994-2011 H. Peter Anvin et al
boot:
  bootloader fazbd backup factoryreset
boot: bootloader_
```

Bootloader Main Page

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. [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/n/c]cancel: 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 Security Event Manager, 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 `fazbdctl` CLI commands on the BigData Controller instead (see [Soft reset FortiAnalyzer-BigData on page 110](#)).



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. 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 `xfstools` 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 90](#) and [Data restore on page 96](#).



You cannot disable this command afterward if it's not needed anymore.

When you back up your data, FortiAnalyzer-BigData backs up the data to an internal HDFS in the Security Event Manager. To back up the data to an external HDFS, all the HDFS nodes must be able to access the external network. By default, all the Security Event Manager hosts (except the Security Event Manager Controller) have no external network access. To allow the rest of the nodes to have external network access, run the following command on the Security Event Manager Controller:

```
fazbdctl enable ip-forward
```

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

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. The upgrade starts with the FortiAnalyzer-BigData main host and then the Security Event Manager hosts. During the upgrade, the GUI is not available. Log collecting, LogView, and FortiView operations are also not available.

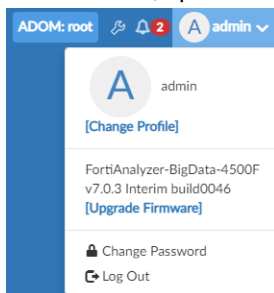


It is recommended to perform the upgrade via the GUI.

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

To upgrade FortiAnalyzer-BigData with the GUI:

1. In the banner, open the Account menu and click *Upgrade Firmware*.




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


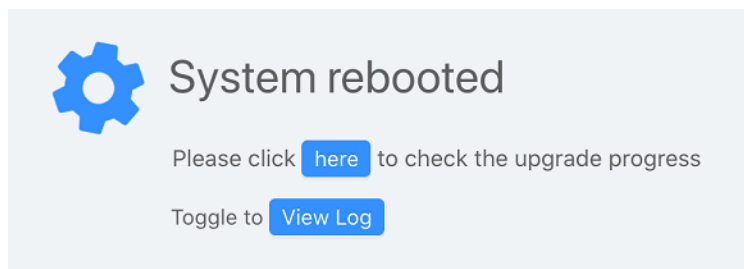
3. Enter the FTP server's IP address, username, password, and file path.

4. 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.

**5.** 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 112](#).

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 [To connect to the Security Event Manager Controller: on page 31](#).



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 fazbd -U <ftp_path> -u <user> -p <password>
```

Or, interactively,

```
fazbdctl upgrade fazbd
```

3. Follow the onscreen instructions to enter your FTP 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%] Finished setting up Core Services.
```



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

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

You can scale out by adding more extender chassis.



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

To add an extender chassis

1. Ensure both the main and extender chassis are running the same version. 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 (see [To connect to the Security Event Manager Controller: on page 31](#)) and run the following command:

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



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. Power off Blade A1 in the extender chassis from CMM.
4. Connect the 40GE links with QSFP between the Internal Switch Modules (Switch Module #1) of the extender and main chassis.

5. 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
```

 There should be 13 additional hosts added as members. Wait until all the hosts' status shows as *Alive*.
6. Access the FortiAnalyzer-BigData GUI of the main chassis, and go to *Cluster Manager > Hosts*.
7. Click *Assign Role* to assign the newly added hosts.
 New hosts should have a "new" label.
8. Wait for the "Assign" job to complete for all services to become healthy.
9. Recommended. After scaling out, rebalance the data across the cluster. See [How to rebalance the data on page 115](#).

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:

1. Access the Security Event Manager Controller of the main chassis (see [To connect to the Security Event Manager Controller: on page 31](#)) and run the following command to stop the DHCP service.

```
systemctl stop dhcpd
```
2. Run the following command to reset all all hosts connected to the current controller:

```
fazbdctl unstack-chassis
```
3. Disconnect the connection between the Internal Switch Modules (Switch Module #1) between the main and extender chassis.
4. Reset and reboot the Bootloader.
 - a. Access the bootloader of the Security Event Manager Controller of the main chassis. See, [Bootloader on page 99](#).
 - b. Enter option 7, `Reset OS and Clear User Data`, to hard reset the host.
 - c. Wait until the Bootloader finishes rebooting.
5. Access the bootloader of Blade A2 on extender chassis. See [Bootloader on page 99](#).
6. Use the `Set Role` command to set the role to `controller`.
7. Access the Security Event Manager Controller of the main chassis and initialize the cluster:

```
fazbdctl init cluster
```
8. Access the Security Event Manager Controller of the extender chassis and initialize the cluster:

```
fazbdctl init cluster
```

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 110](#)
- [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 Security Event Manager hosts while keeping all data volume intact.
- Reset the FortiAnalyzer-BigData power.
- Align all the blade 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. If you do not know the version of your bootloader, it is recommended that you upgrade the bootloader before proceeding.

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 31](#)) 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%] Finished setting up Core Services.`

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 blade 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 reset your FortiAnalyzer-BigData:



Before proceeding with the steps below, your version of FortiAnalyzer-BigData bootloader must match your current version of FortiAnalyzer-BigData. If you do not know the version of your bootloader, it is recommended that you upgrade the bootloader before proceeding.

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 31](#)), 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%] Finished setting up core services.
```

Troubleshooting

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

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 99).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.

From the Security Event Manager Controller, enter the following command to retry upgrading from where it fails:

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

For more information about connecting to the Security Event Manager Controller, see [To connect to the Security Event Manager Controller: on page 31](#).

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.

Error condition	Troubleshooting suggestion
The Security Event Manager Controller cannot boot up after an upgrade and you cannot connect to the Security Event Manager Controller	<p>Perform the following steps:</p> <ol style="list-style-type: none"> 1. Access the bootloader of the Security Event Manager Controller (see Bootloader on page 99). 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 <code>fazbdctl</code> CLI commands (see Upgrade FortiAnalyzer-BigData on page 105) 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 reset command:

- From the Security Event Manager Controller, 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 from the controller to retry soft reset on the cluster from the very beginning:

```
fazbdctl reset cluster -o restart
```

For more information about connecting to the Security Event Manager Controller, see [To connect to the Security Event Manager Controller: on page 31](#).

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 105) and then try resetting again.
The Security Event Manager Controller cannot start or the system is not accessible after a hard reset.	<p>Perform the following steps:</p> <ol style="list-style-type: none"> 1. Access the bootloader of the Security Event Manager Controller (see Bootloader on page 99). 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 <code>fazbdctl</code> 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 99](#)).
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

This section contains instructions on how to gracefully remove and replace a malfunctioning hardware blade running one of the Security Event Manager 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 follows a naming convention: blade-10-0-{chassis ID}-{blade ID}.

A blade named "blade-10-0-1-3" means that "1" represents the chassis ID and the "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 10.0.1.3.

To replace a blade:



Before proceeding with the steps below, your version of FortiAnalyzer-BigData bootloader must match your current version of FortiAnalyzer-BigData. If you do not know the version of your bootloader, it is recommended that you upgrade the bootloader before proceeding.

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 from Security Event Manager Controller: `fazbdctl enable pxe`
3. Insert the replacement blade, and power it on.
4. From the bootloader (see [Bootloader on page 99](#)), set the chassis ID and the blade ID of the replacement blade to match the one from *Cluster Manager > Hosts*.
5. Monitor the status of the blade from CMM (see Remotely control blades via CMM on page 20). Wait until the host OS completes booting. This may take 5-10 minutes.
6. After the host boots up and joins the cluster, it will appear in *Cluster Manager > Hosts* web GUI.
7. From the *Hosts* page, click *Assign Role* to recover the role on the host. The new host will have a *pending* label. When the *Assign Role* job completes, the blade replacement is done.

8. (Recommended) After the *Assign Role* job is complete, rebalance the data across the cluster. See [How to rebalance the data on page 115](#).

How to reset a single host

This section contains instructions on how to gracefully reset a software malfunctioned Security Event Manager 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 follows a naming convention: blade-10-0-{chassis ID}-{blade ID}.

A blade named "blade-10-0-1-3" means that "1" represents the chassis ID and the "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 10.0.1.3.

To reset a single host:

1. Access the bootloader of the malfunctioned host (see [Bootloader on page 99](#)), enter the Reset OS option and wait until it finishes rebooting.
2. Monitor the status of the blade from the DMM (see *Remotely control blades via CMM*) 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.
5. (Recommended) After the *Assign Role* job is completed, rebalance the data across the cluster. See [How to rebalance the data on page 115](#)

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

Choose a maintenance window when the log receiving rate is low.

To rebalance the data:

1. Go to *Cluster Manager > Services > Core* and stop the query, *Ingestion and Pipeline services*.
2. Go to *Cluster Manager > Jobs* and click *Create Custom Job*.
3. Select *Data Rebalance (All)* and *Schedule "Manual"*, then click *Create*.
4. 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.
5. 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*.
6. Go to *Cluster Manager > Services > Core* and resume the stopped services.

How to recover from an unhealthy service status

The service levels in the Security Event Manager is highly available and fault tolerant with data is 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 43](#)). 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 103](#) 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.

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



The Kudu 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 Kudu 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 90](#)). 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 consensus mismatch

In rare situations, the Kudu tablet consensus may break, such as when an ungraceful host is powered off. When this occurs, *Monitor* > *Health* > *Kudu Health Check* will fail and report `CONSENSUS_MISMATCH` in the check result.

Example:

```
Tablet fcdb22e988f54674bf7bd81957d96d99 of table db_log_public.__root_fgt_ipfix 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
```

To fix the consensus mismatch:

1. Go to *Cluster Manager* > *Jobs* > *Create Custom Job*.
2. From the *Template* dropdown, select *Kudu Replica Rebuild* and configure the following settings:
 - *Tablet Server Address*: Use any of the two non-leader replica's hostnames from the table conflict output.
 - *Tablet Id*: Use the *Tablet id* that appears in the first line of the table conflict output.
 - *Reason*: Enter a description of the error.

In the example above, the fields will be configured as follows:

3. In the *Jobs* table view, find the *Kudu Replica Rebuild* row, and click *Run in Actions*.
4. Repeat Steps 1-3 if there are multiple conflicts in the Kudu Health Check result to run one job against each of the conflicted tablets.

After the job is submitted, the tablet goes into `recovering` mode (see the example Kudu Health Check result below). The recovery may take several minutes, depending on the tablet size. Run the Kudu Health Check repeatedly until the health check returns `success`.

Example:

```
Tablet fcdb22e988f54674bf7bd81957d96d99 of table 'db_log_public.__root_fgt_ipfix' 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]
```

How to set up management and external IP addresses using CLI

Use the following CLI commands to set up external management IPs on Security Event Manager Controller and Security Event Manager hosts.

Prerequisite

Access the Security Event Manager Controller. See [To connect to the Security Event Manager Controller: on page 31](#).

Setting up management IP address on the Security Event Manager Controller

To set external IP/mask and gateway information on the Security Event Manager Controller, run the following command.

```
fazbdctl set addr {external IP/mask} [<gateway>] --management
```

To allow the DHCP server to assign external IP/mask on the Security Event Manager Controller, run the following command.

```
fazbdctl set addr dhcp --management
```

Setting up external IP address on a single Security Event Manager host

To explicitly set external IP/mask and gateway information on a Security Event Manager host, run the following command.

```
fazbdctl set addr -H <internal IP> <external IP/mask> [<gateway>]
```

To allow the DHCP server to assign external IP/mask on a Security Event Manager host, run the following command.

```
fazbdctl set addr -H <internal IP> dhcp
```

Setting up external IP addresses on all Security Event Manager hosts

To set external IP/mask and gateway information on all Security Event Manager hosts, run the following command.

```
fazbdctl set addr <external IP/mask> [<gateway>] -A
```



An optional flag is used to set external IP addresses on all Security Event Manager hosts from the Security Event Manager Controller. In this case, the `<external IP/mask>` field specifies the starting external IP address to be assigned to the first Security Event Manager host.

The remaining Security Event Manager hosts are assigned external IP addresses incrementally from the starting external IP address within the network subnet, wrapping around when the boundary of network subnet is reached. This optional flag does not support DHCP.

Example

In the following example cluster configuration:

Role	Address	Ext Address
controller	10.0.1.2	
member	10.0.1.32	
member	10.0.1.33	
member	10.0.1.34	
member	10.0.1.35	
member	10.0.1.36	

After running CLI command

```
fazbdctl set addr 10.106.2.173/24 10.106.2.254 -A
```

The new cluster configuration becomes:

Role	Address	Ext Address
controller	10.0.1.2	10.106.2.173
member	10.0.1.32	10.106.2.174
member	10.0.1.33	10.106.2.175
member	10.0.1.34	10.106.2.176
member	10.0.1.35	10.106.2.177
member	10.0.1.36	10.106.2.178

Clearing external IP addresses on Security Event Manager hosts

To clear external IP/mask and gateway information on a Security Event Manager host, run the following command.

```
fazbdctl unset addr -H <internal IP>
```

To clear external IP/mask information on all Security Event Manager hosts, run the following command.

```
fazbdctl unset addr -A
```

Displaying external IP addresses on Security Event Manager Controller and hosts

To display external IP address information on the Security Event Manager Controller and hosts, run the following command.

```
fazbdctl show members
```

Change Log

Date	Change Description
2021-11-19	Initial release.
2022-01-26	Updated Upgrade FortiAnalyzer-BigData on page 105.
2022-02-04	Updated How to rebalance the data on page 115, Connect to the FortiAnalyzer-BigData CLI on page 31, Connect to the Chassis Management Module on page 16, and Set up the FortiAnalyzer-BigData network on page 14.
2022-02-08	Updated How to scale out on page 108.
2022-11-09	Updated What to do if an upgrade fails on page 112.
2022-11-17	Updated Key terms and concepts on page 7.
2023-01-25	Added Configure LACP on page 28.
2023-02-22	Updated What to do if an upgrade fails on page 112.
2023-03-27	Updated How to replace a blade on page 114.
2023-07-18	Updated How to scale out on page 108.
2023-07-25	Updated How to scale out on page 108.
2023-08-25	Updated Upgrade FortiAnalyzer-BigData on page 105.
2024-01-23	Updated How to scale out on page 108.



www.fortinet.com

Copyright© 2024 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 General Counsel, with a purchaser that expressly warrants that the identified product will perform according to certain expressly-identified performance metrics and, in such event, only the specific performance metrics expressly identified in such binding written contract shall be binding on Fortinet. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. 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.