



# FortiAnalyzer-BigData - Administration Guide

Version 6.2.1

**FORTINET DOCUMENT LIBRARY**

<https://docs.fortinet.com>

**FORTINET VIDEO LIBRARY**

<https://video.fortinet.com>

**FORTINET BLOG**

<https://blog.fortinet.com>

**CUSTOMER SERVICE & SUPPORT**

<https://support.fortinet.com>

**FORTINET TRAINING & CERTIFICATION PROGRAM**

<https://www.fortinet.com/training-certification>

**FORTINET TRAINING INSTITUTE**

<https://training.fortinet.com>

**FORTIGUARD LABS**

<https://www.fortiguard.com>

**END USER LICENSE AGREEMENT**

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

**FEEDBACK**

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



January 23, 2024

FortiAnalyzer-BigData 6.2.1 Administration Guide

58-621-655880-20240123

# TABLE OF CONTENTS

<b>About FortiAnalyzer-BigData</b>	<b>5</b>
Main Features	5
Supported models	6
Key terms and concepts	6
<b>FortiAnalyzer-BigData Hardware environment</b>	<b>8</b>
<b>Set up process</b>	<b>9</b>
Initial set up	9
Set up the FortiAnalyzer-BigData network	10
Set up Administrator accounts	12
Connect to the Chassis Management Module	12
Set up the CMM network	13
Configure the CMM password	16
Configure the Blade Management Network	17
Remotely control blades via CMM	18
Configure the BMC password	18
Connect to the FortiAnalyzer-BigData CLI	22
<b>GUI overview</b>	<b>23</b>
Cluster Manager	23
Custom refresh settings	25
Commands management	25
Notifications management	26
<b>Host management</b>	<b>27</b>
Role assignment	28
<b>Service management</b>	<b>29</b>
Service groups	30
Service details	30
<b>Monitor</b>	<b>33</b>
Dashboards	33
Filtering the Dashboard	34
Customizing the Dashboard	34
Logs and metrics	36
Explore logs	37
Explore metrics	41
Health	43
Health Check	43
Alert	44
<b>Job management and automation</b>	<b>48</b>
Job history	49
Built-in automation jobs	50
Custom automation jobs	50
Custom job templates	52

<b>Data management</b>	<b>56</b>
Manage data lifecycle	57
Data backup	57
Incremental backups	59
Incremental backups	60
Data restore	62
<b>Bootloader</b>	<b>65</b>
Bootloader Main Page	65
1. Configure Network	66
2. Install OS	66
3. Set Role	67
4. Set Chassis ID	67
5. Set Blade ID	67
6. Reset OS	67
7. Reset OS and Clear User Data	68
8. Upgrade Bootloader	68
0. Reboot	68
sh. shell	69
<b>General maintenance and best practices</b>	<b>70</b>
Backup and restore to external HDFS	70
Schedule maintenance tasks for off-peak hours	70
Maintain database integrity	71
<b>Upgrade FortiAnalyzer-BigData</b>	<b>72</b>
<b>Scaling FortiAnalyzer-BigData</b>	<b>75</b>
How to scale out	75
How to remove a chassis from a stacked setup	76
Remove an extender chassis	76
<b>Reset FortiAnalyzer-BigData</b>	<b>77</b>
Soft reset FortiAnalyzer-BigData	77
Hard reset FortiAnalyzer-BigData	77
<b>Troubleshooting</b>	<b>79</b>
What to do if an upgrade fails	79
What to do if a soft reset fails	79
What to do if a hard reset fails	80
How to repair disk failures	80
How to replace a blade	81
How to reset a single host	82
How to recover from an unhealthy service status	83
Core services	84
Data Lake services	84
Message Broker services	85
How to recover from a full disk	86
<b>Change Log</b>	<b>87</b>

# 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 6.2.1. For a list of supported FortiGate models, see the [FortiAnalyzer 6.2.1 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 to perform 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.

### Security Event Manager Host(s)

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.

### Controller

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

### 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 GUI 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 75](#)), 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
  - HDFS Datanode
  - Kafka Broker
  - Kudu Master
  - Kudu Tablet Server
  - Yarn Node Manager
  - Zookeeper
- MetaStore Node
  - HDFS Datanode
  - HDFS Namenode
  - Kafka Broker
  - Kudu Tablet Server
  - Yarn Node Manager
  - Yarn Resource Manager
- Data Node
  - HDFS Datanode
  - Impala
  - Kafka Broker
  - Kudu Tablet Server
  - Yarn Node Manager

## Services

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

# 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 the GUI.
- The remaining 13 blades, also known as the Security Event Manager hosts, are responsible for log storage 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 12](#).



# 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 9](#)
2. [Set up the FortiAnalyzer-BigData network on page 10](#)
3. [Set up Administrator accounts on page 12](#)

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:

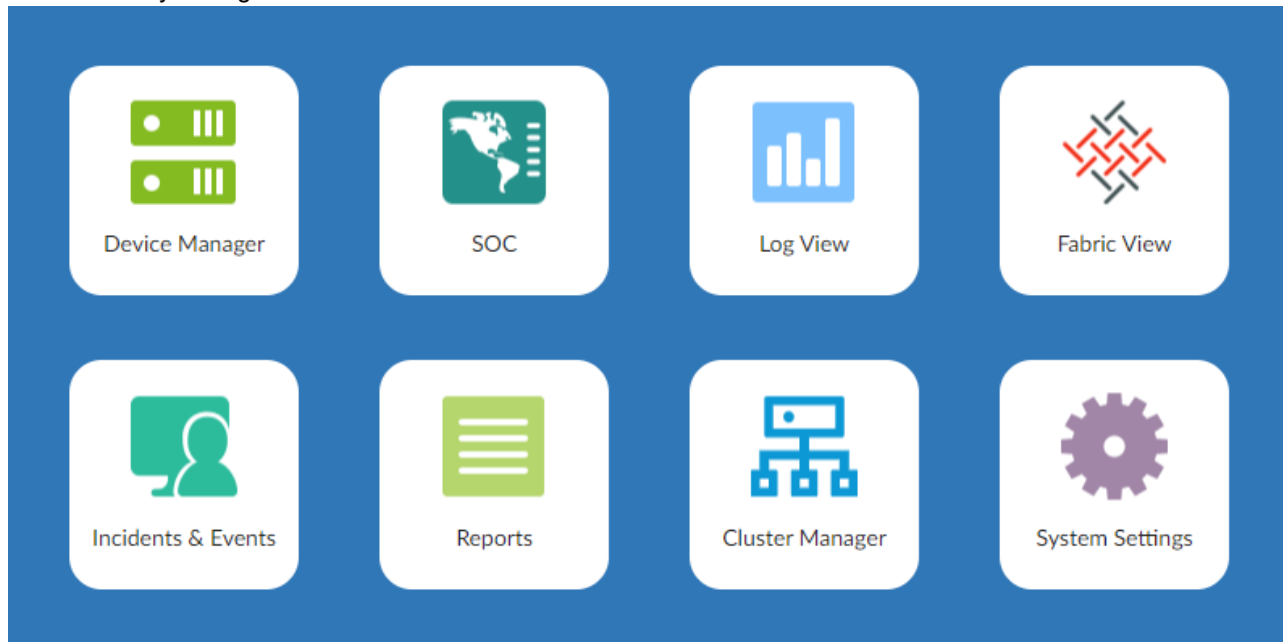
- Ethernet cable
- SPF RJ45 transceiver module
- Management computer

## 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://192.168.1.99>.

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.

### To set up the FortiAnalyzer-BigData network:

1. From the FortiAnalyzer-BigData GUI, go to *System Settings > Network*.
2. Change the *IP Address/Netmask* field 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 etc.
3. Keep the default *Administrative Access* settings.
4. Specify a *Default Gateway*.

5. Click *Cluster Management* to configure the network for the FortiAnalyzer-BigData Cluster Manager.

### System Network Management Interface

Name	port2
IP Address/Netmask	<input type="text" value="10.3.112.95/255.255.255.0"/>
IPv6 Address	<input "::="" 0"="" type="text" value=""/>
Administrative Access	<input checked="" type="checkbox"/> HTTPS <input checked="" type="checkbox"/> HTTP <input checked="" type="checkbox"/> PING <input checked="" type="checkbox"/> SSH <input checked="" type="checkbox"/> SNMP <input checked="" 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
Default Gateway	<input type="text" value="172.16.96.1"/>
Primary DNS Server	<input type="text" value="172.16.100.100"/>
Secondary DNS Server	<input type="text" value="172.16.100.80"/>

[All Interfaces](#)[Routing Table](#)[IPv6 Routing Table](#)[Cluster Management](#)[Apply](#)

- a. In the Cluster Management page, change the *IP Address/Netmask* field to your internal network. This field is different from the one used in step 2. This IP Address is for the FortiAnalyzer-BigData Security Event Manager Controller.
- b. Configure the *Gateway* field as needed. This address is usually the same as the *Default Gateway* field in step 4.

### Cluster Management

IP Address/Netmask	<input type="text" value="10.3.112.96/255.255.255.0"/>
Gateway	<input type="text" value="172.16.96.1"/>

[OK](#)[Cancel](#)

- c. Click *OK* to save your changes.
6. Click *Apply* to save your change.
7. From your management computer, change the IP Address/Netmask accordingly to reconnect it to FortiAnalyzer-BigData.

## 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.

## 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 16](#).

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*.

**IPv4 Setting**

IP Address	172.31.34.169
Subnet Mask	255.255.0.0
Gateway	172.31.0.1
DNS Server IP	10.2.1.205

**IPv4 Setting when DHCP fails**

IP Address	192.168.100.100
Subnet Mask	255.255.255.0
Gateway	192.168.100.1

**IPv6 Setting**

IPv6 Address

☒ Add IP ☐ Delete IP ☒ Auto Configuration

☒ DHCPv6 Stateless ☐ DHCPv6 Stateful

Address List

DNS Server IP

DUID

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 10.160.81.11
SET NETMASK 255.255.255.0
SET GATEWAY 10.160.81.1
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

### To configure the Blade Management Network:

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.

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

IPv4 Setting

IP Scale	<input type="text" value="4"/>
Base IP Address	<input type="text" value="000.000.000.000"/>
Subnet Mask	<input type="text" value="000.000.000.000"/>
Gateway	<input type="text" value="000.000.000.000"/>
DNS Server IP	<input type="text" value="000.000.000.000"/>
VLAN ID	<input type="text" value="0"/>

☐ Apply above setting to all blades and EFFECTIVE all the time. (always autoly apply to the Blades which are re-plugged in)

Save

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

3. Select how you want to obtain an IP address.
  - **Obtain an IP address automatically:** Obtain an IP address automatically by using DHCP.
  - **Use the following IP address:** Set up the IP address by entering the information in the *IPv4 Settings* fields. 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.
    - i. Enter your IP information.
    - ii. In the *IP Scale* field, select a number so that each blade IP address increases by a base of 1, 2, or 4.
4. Check the last box if you want to apply the network setting to all blades. This preserves the Blade IPMI network setting whenever a blade is re-connected.
5. Click **Save**.

## 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 65](#)).

---

## 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:bab2/64	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**

## 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 Cluster Management IP you configured in [Initial set up on page 9](#).
2. Log in using the administrator credentials you created in [Set up Administrator accounts on page 12](#).  
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 9](#).
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](#).

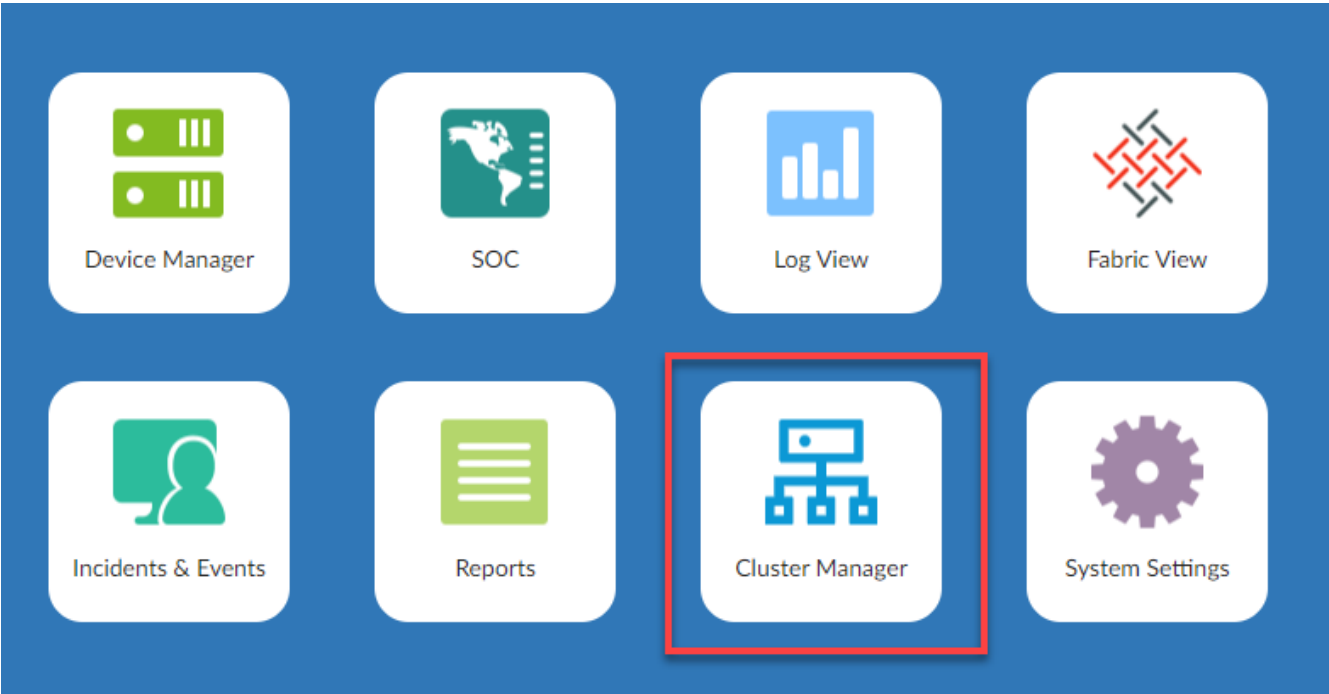


Fortinet strongly recommends that you update the password with the `passwd` command.

---

# GUI overview

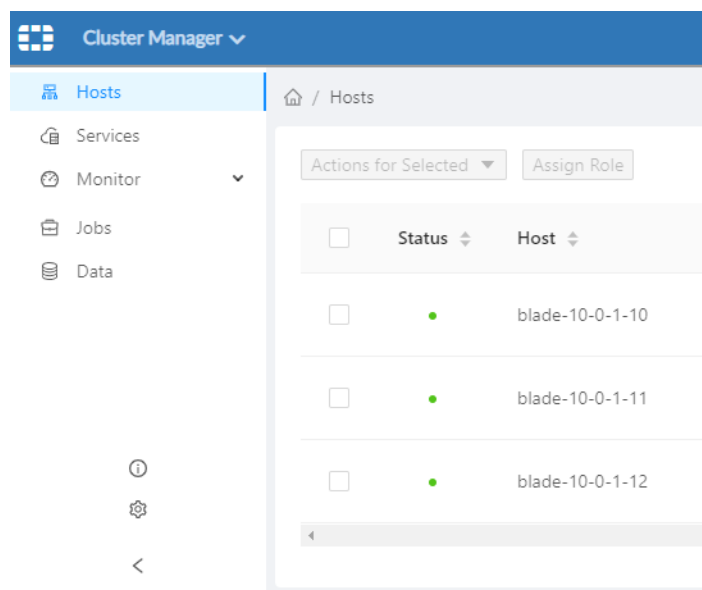
FortiAnalyzer-BigData retains the same general GUI as the base FortiAnalyzer, however, there is a new Cluster Manager tile that enables you to manage all the resources related to the Security Event Manager.





Cluster Manager	The Cluster Manager module enables you to manage hosts, services, logs, queries, jobs, and data resources in the Security Event Manager. See <a href="#">Cluster Manager on page 23</a> .
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 <a href="#">FortiAnalyzer administration guide</a> .

## 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
<b>Hosts</b>	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 <a href="#">Host management on page 27</a> .
<b>Services</b>	The Services page enables you to manage the configurations and life cycle of the Security Event Manager. For more information, see <a href="#">Service management on page 29</a> .
<b>Monitor</b>	The Monitor section contains three pages: <ul style="list-style-type: none"> <li>• <b>Dashboard:</b> Provides a customizable visualization for system metrics.</li> <li>• <b>Log and Metrics:</b> Contains an Explorer tool that enables you to search the logs and metrics that FortiAnalyzer-BigData produces.</li> <li>• <b>Health:</b> Provides push notifications for system health checks and other events.</li> </ul> For more information, see <a href="#">Monitor on page 33</a> .
<b>Jobs</b>	The Jobs page manages system jobs and custom jobs. <ul style="list-style-type: none"> <li>• 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.</li> <li>• Custom jobs can be set up with built-in templates or customizable playbooks.</li> </ul> For more information, see <a href="#">Job management and automation on page 48</a> .
<b>Data</b>	The Data page enables you to manages the data life cycle of your storage groups as well as data backups and restores. For more information, see <a href="#">Data management on page 56</a> .
<b>System Information</b> 	Click to see the current system version number.
<b>System Upgrade</b> 	Click to see your current system version and to upgrade FortiAnalyzer-BigData.




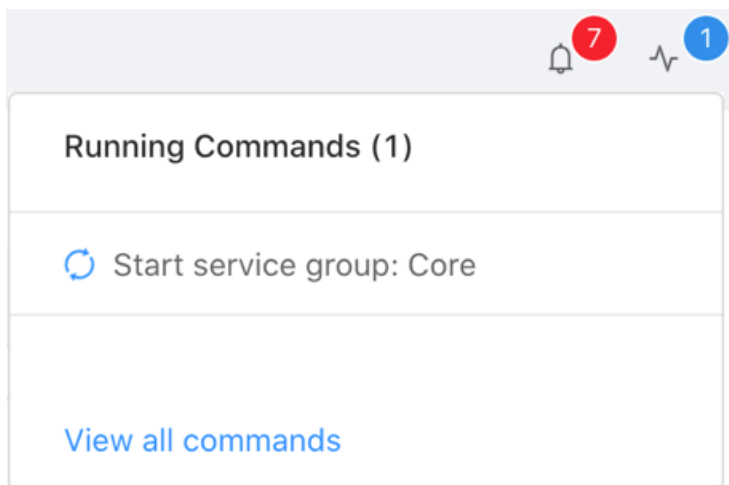
## Custom refresh settings

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.

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

The screenshot shows the 'Commands' page. At the top, there's a 'Running Commands' section with a single entry: 'Start service group: Core'. Below this is the 'Recent Commands' section, which contains a table of command history.

Status	Command Name	Start Time	Duration
✓	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

At the bottom right of the table, there is a pagination control showing '1' selected, with '2' and '3' as options, and navigation arrows.

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

The legend shows two icons with their corresponding meanings:

- 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.

The screenshot shows the Notifications dropdown menu. It has a title 'Unhealth Alerts' and a refresh icon. Below the title, there are two items:

- Unhealth Services** with a count of 5.
- Unhealth Health Check** with a count of 2 and a 'Refresh' button.

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)				<a href="#">Restart</a> <a href="#">Status Details</a>
<input type="checkbox"/>		blade-10-0-1-32	Master Node	10.0.1.32	▶ 7 Instance(s)				<a href="#">Restart</a> <a href="#">Status Details</a>
<input type="checkbox"/>		blade-10-0-1-33	MetaStore Node	10.0.1.33	▶ 6 Instance(s)				<a href="#">Restart</a> <a href="#">Status Details</a>
<input type="checkbox"/>		blade-10-0-1-34	MetaStore Node	10.0.1.34	▶ 9 Instance(s)				<a href="#">Restart</a> <a href="#">Status Details</a>
<input type="checkbox"/>		blade-10-0-1-35	Data Node	10.0.1.35	▶ 5 Instance(s)				<a href="#">Restart</a> <a href="#">Status Details</a>
<input type="checkbox"/>		blade-10-0-1-36	Data Node	10.0.1.36	▶ 5 Instance(s)				<a href="#">Restart</a> <a href="#">Status Details</a>
<input type="checkbox"/>		blade-10-0-1-37	Data Node	10.0.1.37	▶ 5 Instance(s)				<a href="#">Restart</a> <a href="#">Status Details</a>

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"><li> The host is healthy.</li><li> The host is in poor health.</li><li> A command is currently running on the host.</li></ul>
Host Name	The name of the host.
Role Type	Each host can have one of four roles. For more information about each role, see <a href="#">Roles on page 7</a> . <ul style="list-style-type: none"><li>• <b>Master Node</b></li><li>• <b>MetaStore Node</b></li><li>• <b>Data Node</b></li><li>• <b>Unassigned:</b> The host is new and does not have an assigned role. Click <i>new</i> to assign a role to that host (see <a href="#">Role assignment on page 28</a>).</li></ul>
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
<b>CPU Usage</b>	The percentage of the CPU being used.
<b>Memory Usage</b>	How much memory is being used.
<b>Disk Usage</b>	How much space is being used on a disk.
<b>Actions</b>	<p>You can perform the following actions on each host:</p> <ul style="list-style-type: none"> <li>• <b>Restart:</b> Restart the host.</li> <li>• <b>Status Details:</b> See the full metrics view of the host.</li> <li>• <b>Assign Role:</b> Assign a role to a new host.</li> </ul>

## 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)	<a href="#">Assign Role</a>

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.

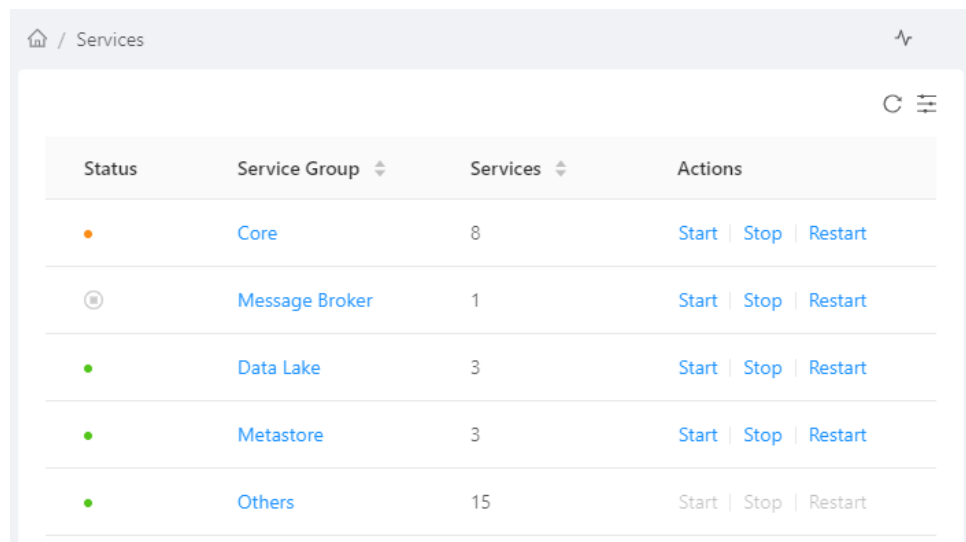


At this time, you can only assign the Data Node role.

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. At the top, there is a breadcrumb 'Home / Services' and a refresh icon. Below the header, there is a table with the following data:

Status	Service Group	Services	Actions
	Core	8	<a href="#">Start</a>   <a href="#">Stop</a>   <a href="#">Restart</a>
	Message Broker	1	<a href="#">Start</a>   <a href="#">Stop</a>   <a href="#">Restart</a>
	Data Lake	3	<a href="#">Start</a>   <a href="#">Stop</a>   <a href="#">Restart</a>
	Metastore	3	<a href="#">Start</a>   <a href="#">Stop</a>   <a href="#">Restart</a>
	Others	15	<a href="#">Start</a>   <a href="#">Stop</a>   <a href="#">Restart</a>

The Services table contains the following columns:

Column header	Description
<b>Status</b>	There are five icons that represent the status of the host: <ul style="list-style-type: none"><li> The services are healthy.</li><li> A command is currently running.</li><li> There is a problem with the service.</li><li> Services within the service group are experiencing issues.</li><li> The service has stopped.</li></ul>
<b>Service Group</b>	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 <a href="#">Service groups on page 30</a> ).
<b>Services</b>	The number of services running in each group.
<b>Actions</b>	There are three actions you can perform on each Service Group or service. <ul style="list-style-type: none"><li>• <b>Start:</b> Start the service group or a specific service.</li><li>• <b>Stop:</b> Stop the service group or a specific service.</li><li>• <b>Restart:</b> Restart the service group or a specific service.</li></ul>

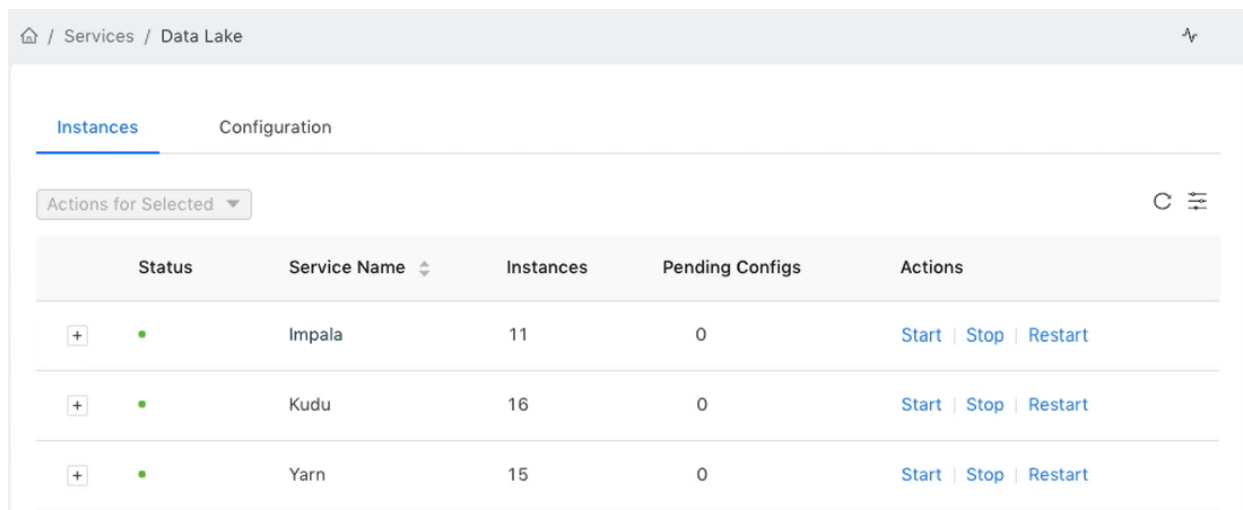
## Service groups

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
<b>Core</b>	<ul style="list-style-type: none"> <li>• Catalog</li> <li>• Query</li> <li>• Ingestion</li> <li>• Data Explorer</li> <li>• Pipeline</li> <li>• Monitor</li> </ul>
<b>Message Broker</b>	<ul style="list-style-type: none"> <li>• Kafka</li> </ul>
<b>Data Lake</b>	<ul style="list-style-type: none"> <li>• Impala</li> <li>• Kudu</li> <li>• Yarn</li> </ul>
<b>Metastore</b>	<ul style="list-style-type: none"> <li>• Zookeeper</li> <li>• HDFS</li> </ul>

## Service details

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



<a href="#">Home</a> / <a href="#">Services</a> / <a href="#">Data Lake</a>				
<a href="#">Instances</a> <a href="#">Configuration</a>				
Actions for Selected <span>⌵</span>				
Status	Service Name	Instances	Pending Configs	Actions
<div>+</div> <div>●</div>	Impala	11	0	<a href="#">Start</a>   <a href="#">Stop</a>   <a href="#">Restart</a>
<div>+</div> <div>●</div>	Kudu	16	0	<a href="#">Start</a>   <a href="#">Stop</a>   <a href="#">Restart</a>
<div>+</div> <div>●</div>	Yarn	15	0	<a href="#">Start</a>   <a href="#">Stop</a>   <a href="#">Restart</a>

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
	Kafka	8	0	<a href="#">Start</a>   <a href="#">Stop</a>   <a href="#">Restart</a>

<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	<a href="#">Start</a>   <a href="#">Stop</a>   <a href="#">Restart</a>
<input type="checkbox"/>		Kafka Broker	Started	blade-10-0-2-32	10.0.2.32	<a href="#">Start</a>   <a href="#">Stop</a>   <a href="#">Restart</a>
<input type="checkbox"/>		Kafka Broker	Started	blade-10-0-2-33	10.0.2.33	<a href="#">Start</a>   <a href="#">Stop</a>   <a href="#">Restart</a>

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.

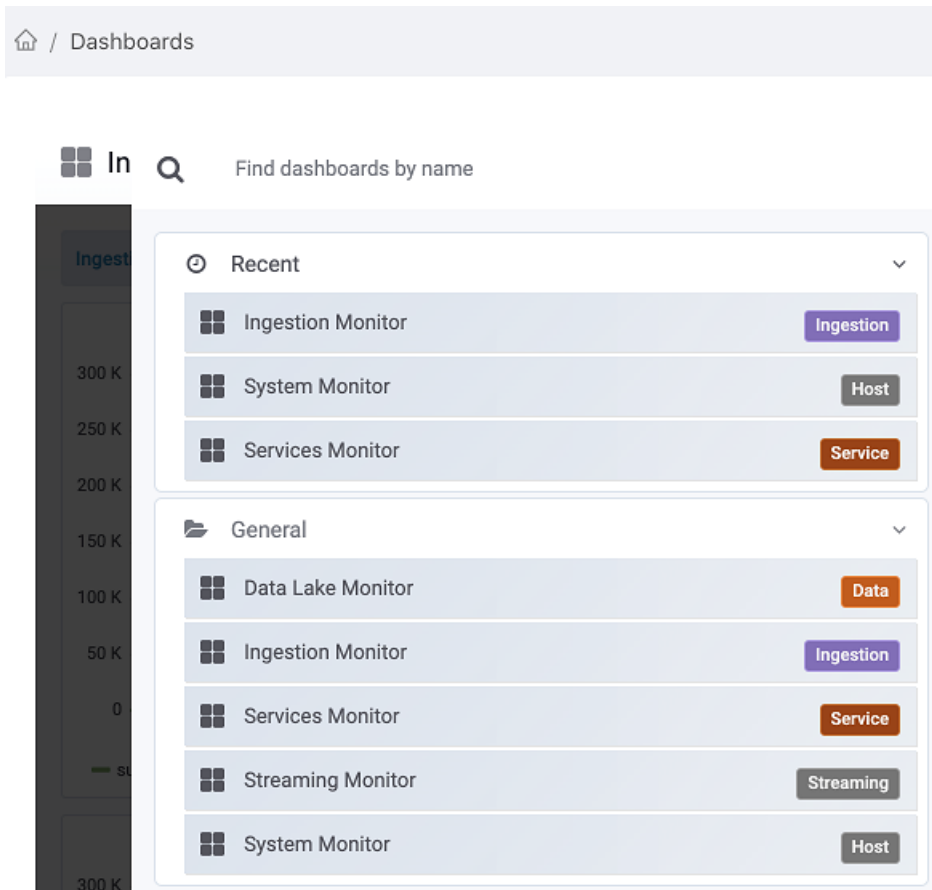
From the Dashboard, you can:

- 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
<b>Add panel</b>	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"> <li>• Add Query: Choose what metrics to track,</li> <li>• Choose Visualization: Choose how you want to visualize the data.</li> <li>• Convert to row: Convert a group of panels into a collapsible row.</li> </ul>
<b>Share Dashboard</b>	Share the dashboard with a link or by exporting a JSON file.
<b>Save Dashboard</b>	Save all the changes you've made to the dashboard.
<b>Dashboard settings</b>	
<b>General</b>	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 <a href="#">Grafana documentation</a> .
<b>Annotations</b>	Add annotations to mark points on a graph.
<b>Variables</b>	Add variables to change the data being displayed in the dashboard.
<b>Links</b>	Add a link to your dashboard so you can go to other dashboards and websites directly.
<b>Versions</b>	See the revision version history for the dashboard.
<b>JSON Model</b>	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
<b>View</b>	Enlarge the panel to see a more detailed view of the graph.
<b>Edit</b>	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.
<b>Share</b>	There are two ways to share your panels: <ul style="list-style-type: none"> <li>• Create a direct link to the particular panel.</li> <li>• Create a snapshot of the panel with sensitive data stripped out.</li> </ul>
<b>Explore</b>	View the historical logs and metrics for the panel.
<b>More</b>	
<b>Duplicate</b>	Add a duplicate of the panel to your dashboard.
<b>Copy</b>	Create a copy of the pane. You can paste the panel to the Dashboard from <i>Add panel</i> .
<b>Panel JSON</b>	See the JSON model that defines the panel.
<b>Export CSV</b>	Export a CSV file with panel data.
<b>Toggle Legend</b>	Click to display or conceal the panel legend.
<b>Remove</b>	Remove the panel from the Dashboard.

## Logs and metrics

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

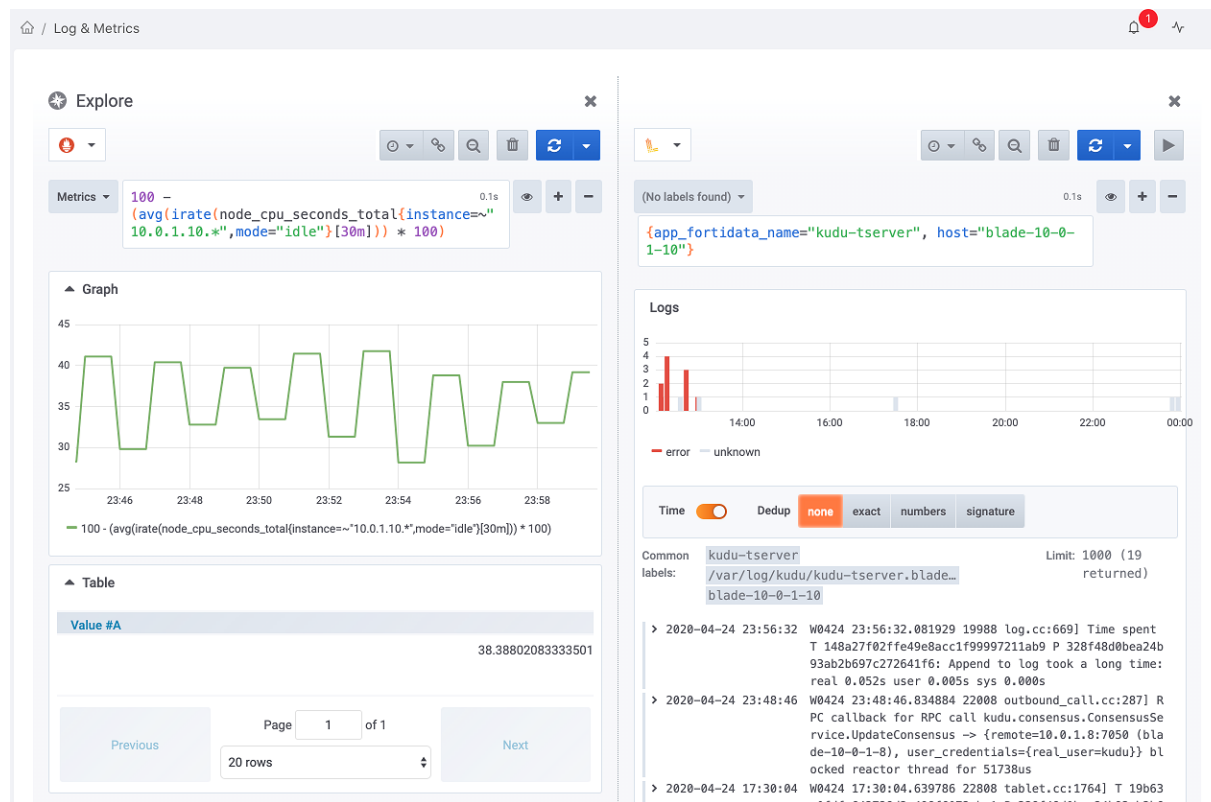
You can use the Explore search tool to switch between the Logs or Metrics view. The default selection is Logs.



- 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.

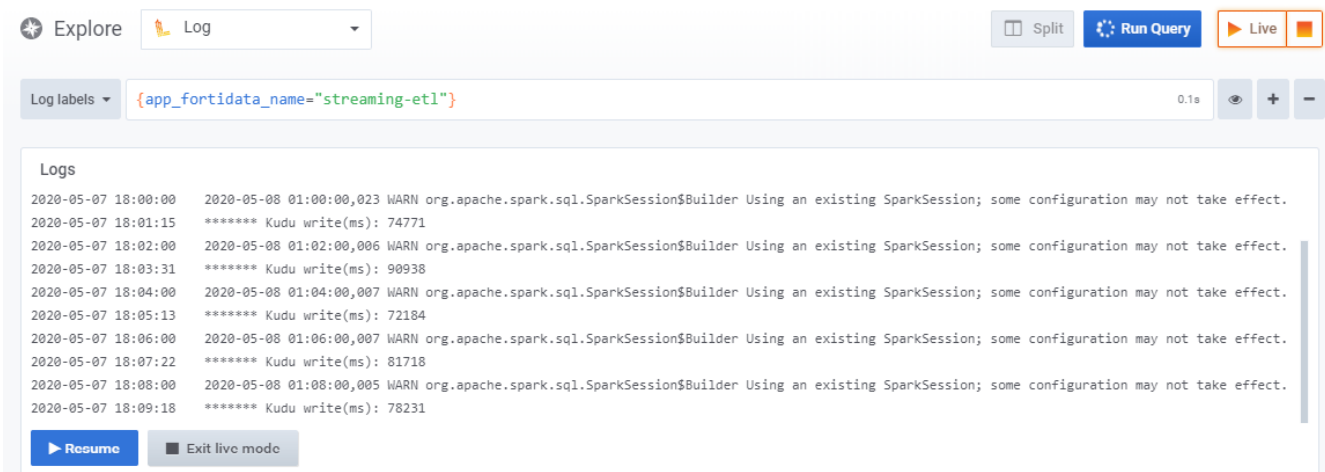
For more details, refer to the Loki query language (LogQL) documentation.

## 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 includes a timestamp, a log entry, and a 'Show context' link. A red box highlights the 'Show context' link for the first result.

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

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

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

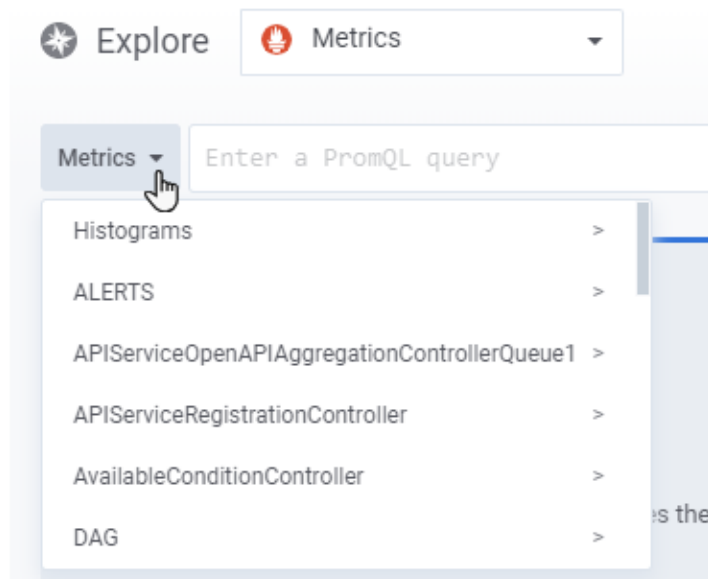
- Found 10 rows. Load 10 more
- get process information of datanode ----- 0.63s
- get available space for each data dir ----- 0.73s
- get data dir content ----- 1.24s
- run command to get datanode usage info ----- 2.56s
- Gathering Facts ----- 3.48s
- =====
- Saturday 25 April 2020 07:49:36 +0000 (0:00:00.030) 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



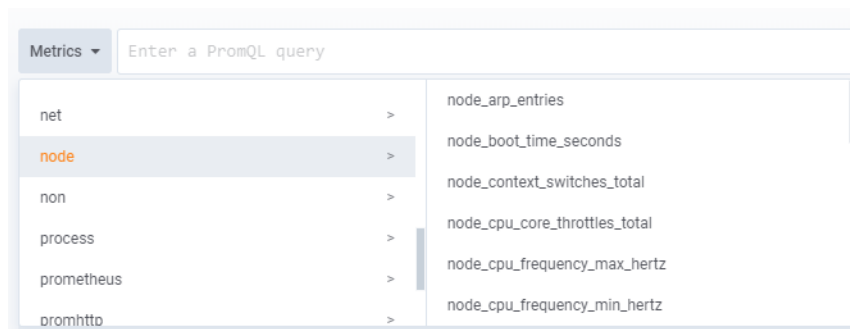
## Explore metrics

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.

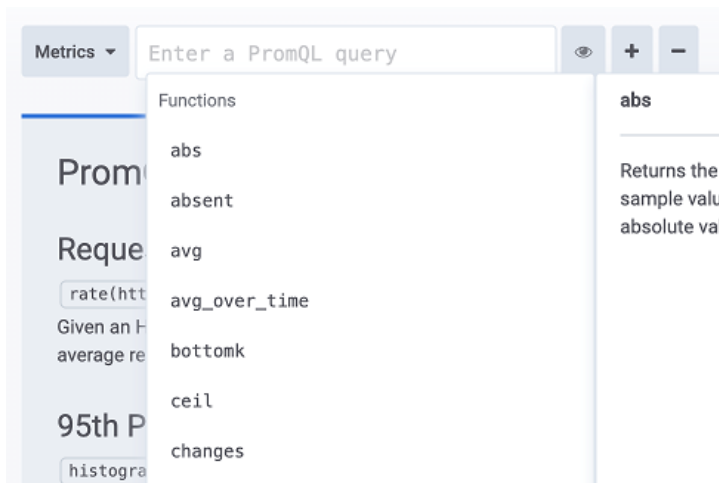


Metrics are grouped by prefixes, for example, all Node metrics are grouped under the "node" prefix.



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.

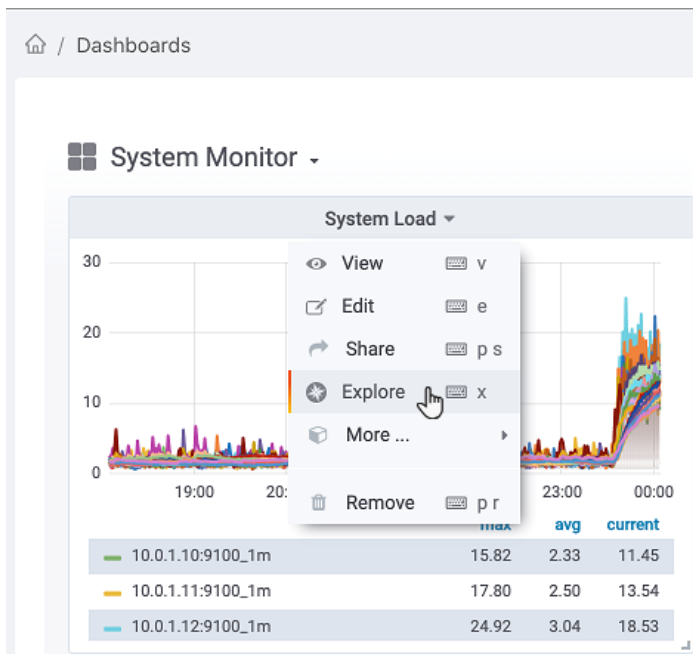


For more details, refer to the [Prometheus Query Language documentation](#).

## 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*.

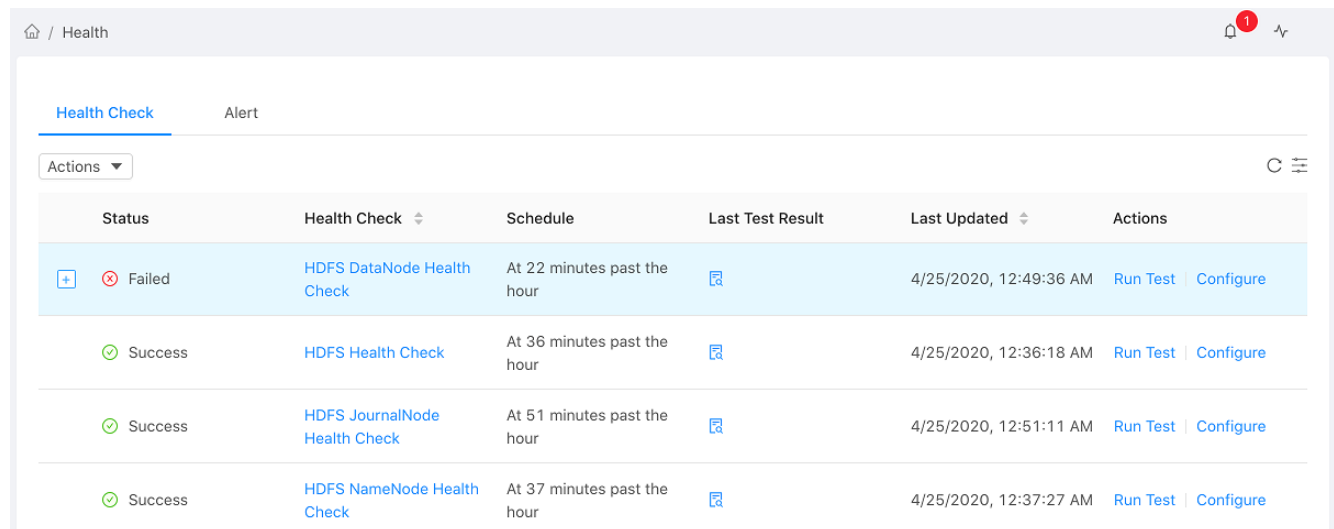


## Health

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 health checks: 'HDFS DataNode Health Check' (Failed), 'HDFS Health Check' (Success), 'HDFS JournalNode Health Check' (Success), and 'HDFS NameNode Health Check' (Success). Each row includes a status icon, the check name, the schedule, the last test result (with a link to view details), the last updated timestamp, and actions ('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	<a href="#">Run Test</a> <a href="#">Configure</a>
Success	HDFS Health Check	At 36 minutes past the hour		4/25/2020, 12:36:18 AM	<a href="#">Run Test</a> <a href="#">Configure</a>
Success	HDFS JournalNode Health Check	At 51 minutes past the hour		4/25/2020, 12:51:11 AM	<a href="#">Run Test</a> <a href="#">Configure</a>
Success	HDFS NameNode Health Check	At 37 minutes past the hour		4/25/2020, 12:37:27 AM	<a href="#">Run Test</a> <a href="#">Configure</a>

The Health Check table contains the following columns:

Column header	Description
<b>Status</b>	<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>
<b>Health Test</b>	<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/>
<b>Schedule</b>	Shows how often the health check is run.
<b>Last Test Result</b>	View the full health test result by clicking <i>Test Result</i> .
<b>Last Updated</b>	The last time the health test was run.
<b>Actions</b>	<p>You can perform two actions on the health test:</p> <ul style="list-style-type: none"> <li><b>Run Test:</b> Manually start the health test.</li> <li><b>Configure:</b> Change how often the test is run by configuring the scheduling settings.</li> </ul>

## 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

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**.

Health Check **Alert**

### Edit Notification Channel

Name	Slack
Type	Slack
Default (send on all alerts)	<input type="checkbox"/>
Include image	<input checked="" type="checkbox"/>
Disable Resolve Message	<input type="checkbox"/>
Send reminders	<input type="checkbox"/>

### Slack settings

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

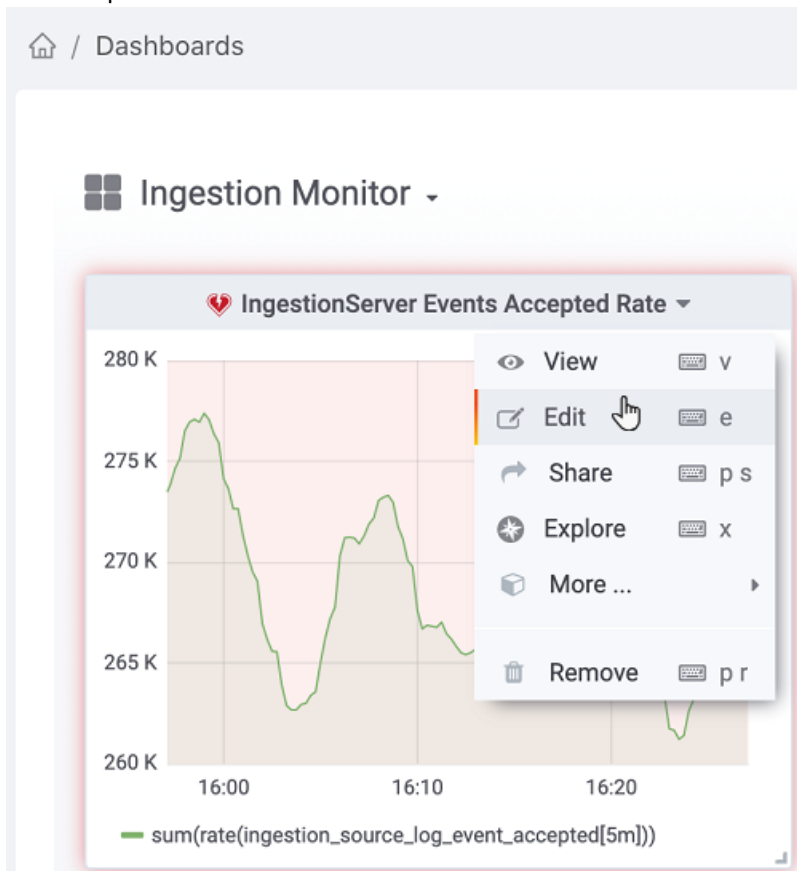
## Custom alert rules

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.

- 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



# Job management and automation

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

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

The Jobs table contains the following columns:

Column header	Description
<b>Summary</b>	The name or short description of a job. You can click the summary to view its execution history (see <a href="#">Job history on page 49</a> ).
<b>Job Type</b>	There are two types of jobs: <ul style="list-style-type: none"> <li><b>Built-in:</b> Pre-configured system jobs.</li> <li><b>Custom:</b> Job created by an administrator.</li> </ul>
<b>Schedule</b>	Shows how often the job is run.
<b>Last Job Status</b>	Indicates the status of the job: <ul style="list-style-type: none"> <li><b>Success:</b> The job execution successful.</li> <li><b>Failed:</b> The job execution failed.</li> <li><b>Running:</b> The job is currently executing.</li> <li><b>Queued:</b> The job has been put into an execution queue and will be executed shortly.</li> <li><b>Timeout:</b> The job execution has timed out.</li> <li><b>Aborted:</b> The job execution has been interrupted. This status usually occurs when the user manually aborts.</li> <li><b>Skipped:</b> 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.</li> </ul>
<b>Last Job Result</b>	View the last job execution result by clicking <b>Job Result</b> <a href="#">View</a> .



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

## Job history

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

🏠 / Jobs / Data Retention
🔔 🔍

Run Job
🔄 ☰

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

<
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:

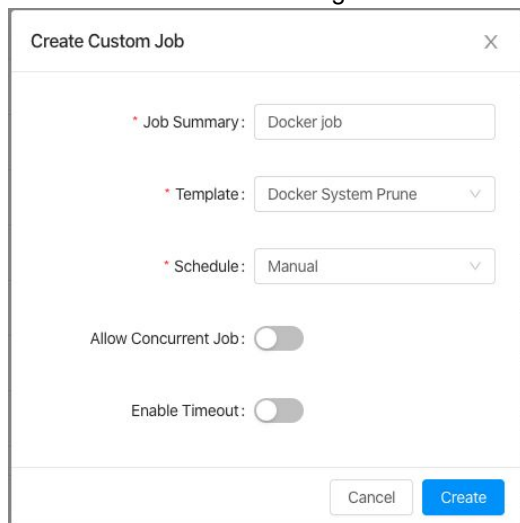
Built-in jobs	Description
<b>Data Retention</b>	Created automatically when storage groups are created. This job is used to apply data retention policies for the Storage Group which marks the old data for deletion and makes space for future data.
<b>Data Rebalance</b>	Created automatically when storage groups are created. This job is used to rebalance Kudu data partitions to evenly distribute them across the Security Event Manager hosts.
<b>Data Appendix</b>	Created automatically when storage groups are created. This job generates the list of available sub-types of FortiGate Event logs for LogView.
<b>Facet Formation - Report</b>	Created automatically when storage groups are created. This job generates the pre-aggregated facets to speed up FortiView queries.
<b>Facet Formation - FortiView</b>	Created automatically when storage groups are created. This job generates the pre-aggregated facets to speed up Report queries.
<b>Storage Group Restore</b>	This job will be created automatically when you launch the storage group restore function from the Data page. For more details, see <a href="#">Data restore on page 62</a> .

## 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.




- Complete the following fields:

Field name	Description
<b>Job Summary</b>	Enter the job description.
<b>Template</b>	Select a job template. For templates that have additional fields to fill out, see <a href="#">Custom job templates on page 52</a> .
<b>Schedule</b>	Select a scheduling timer: <ul style="list-style-type: none"> <li><b>Manual:</b> The job will not be executed until you manually launch it.</li> <li><b>Daily:</b> The job is scheduled to run on a daily basis. Select a run time and enable the <b>Enable Job</b> toggle so the schedule takes effect. To pause the job schedule, disable the toggle.</li> <li><b>Advanced:</b> 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 <b>Enable Job</b> toggle to enable so the schedule takes effect. To pause the job scheduling, disable the toggle.</li> </ul>
<b>Allow Concurrent Job</b>	Enable to allow multiple jobs to run at the same time.
<b>Enable Timeout</b>	Enable to define job timeout.

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

### To import custom jobs:

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



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

### To export multiple custom jobs:

- From the Jobs page, select the jobs you want to export.
- In the top-left corner of the Jobs page, click *Export Job*. The Confirm Export Job dialog box loads.
- 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 group 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"> <li>• hdfs_datanode</li> <li>• hdfs_namenode</li> <li>• kudu_tserver</li> </ul> | These inventory groups can be used to select the host(s) that have the named services running. |
|--|--|

<ul style="list-style-type: none"> <li>• kudu_hive_metastore</li> <li>• zookeeper</li> <li>• kafka_broker</li> <li>• impala_catalog</li> <li>• impala</li> <li>• impala_statestore</li> <li>• yarn_nodemanager</li> <li>• yarn_resource_manager</li> <li>• spark_history_server</li> </ul>	For example, using “ <b>host:</b> kudu_tserver” in your playbook allows it to be executed on all hosts has kudu-tserver instance.
<ul style="list-style-type: none"> <li>• hdfs_datanode_reachable</li> <li>• hdfs_namenode_reachable</li> <li>• kudu_tserver_reachable</li> <li>• kudu_reachable</li> <li>• hive_metastore_reachable</li> <li>• zookeeper_reachable</li> <li>• kafka_broker_reachable</li> <li>• impala_catalog_reachable</li> <li>• impala_reachable</li> <li>• impala_statestore_reachable</li> <li>• yarn_nodemanager_reachable</li> <li>• yarn_resource_manager_reachable</li> <li>• spark_history_server_reachable</li> </ul>	<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 “<b>hosts:</b>kudu_reachable” will randomly return one that is reachable at the execution time.</p>
<ul style="list-style-type: none"> <li>• metastore</li> <li>• datanode</li> <li>• master</li> </ul>	These groups can be used to select hosts the belong to the named role.
<ul style="list-style-type: none"> <li>• metastore_reachable</li> <li>• datanode_reachable</li> <li>• master_reachable</li> </ul>	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"> <li>• controllerlp</li> </ul>	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:

First, select a storage group, 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 83](#)).

### 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 75](#).

### 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 83](#)).

### Purge Data Pipeline

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

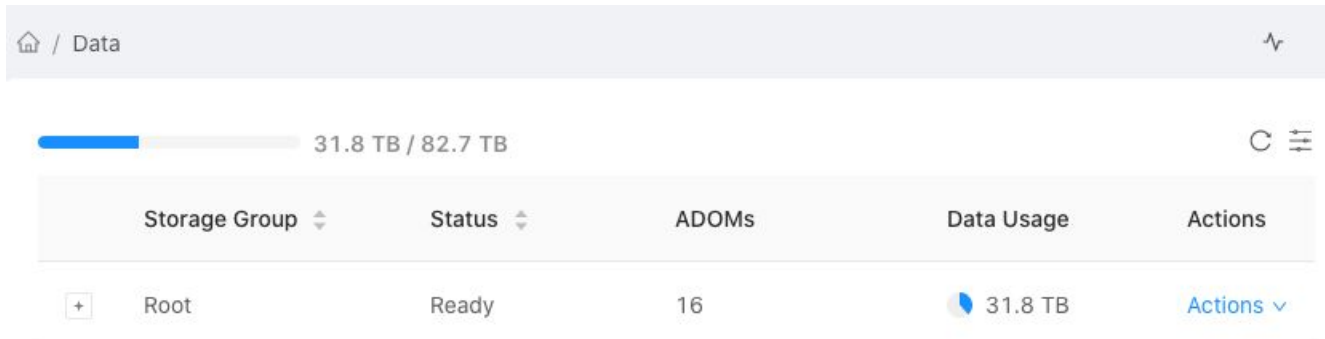


Any unprocessed data will be lost (see [How to recover from an unhealthy service status on page 83](#)).

---

# Data management

FortiAnalyzer-BigData manages the disk space via Storage Groups and the Data page contains a table listing all Storage Groups. By default, FortiAnalyzer-BigData is shipped with default "Root" storage group.



The Storage Group table contains the following columns:

Column header	Description
Storage Group	The name of storage group. You can expand each storage group to display all the ADOMs in that group.
Status	Indicates the status of the storage group. <ul style="list-style-type: none"><li>• <b>Ready</b>: The storage group is ready for use.</li><li>• <b>In Progress</b>: The storage group is being created and is not yet ready for use.</li><li>• <b>Failed</b>: The storage group creation failed.</li></ul>
ADOMs	The number of ADOMs in that storage group.
Data Usage	How much data is in use.
Actions	You can perform the following actions on a storage group: <ul style="list-style-type: none"><li>• <b>Manage Data Lifecycle</b>: Determine how long you want to store the data, and when to do a data rollover. For more information, see <a href="#">Manage data lifecycle on page 57</a>.</li><li>• <b>Manage Job</b>: Manage jobs in that storage group.</li><li>• <b>Backup</b>: Create a backup of that storage group.</li><li>• <b>Restore</b>: Restore data.</li></ul>

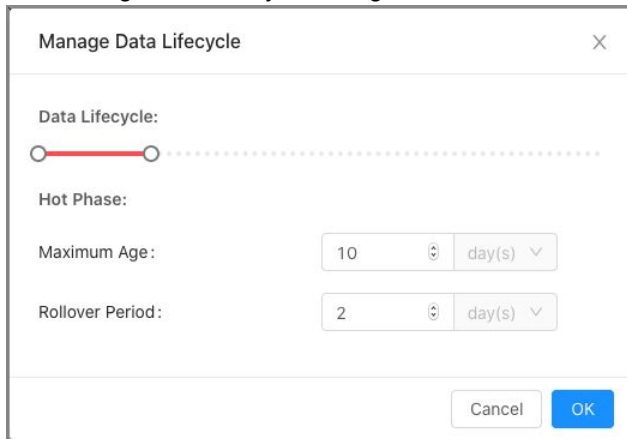


## Manage data lifecycle

You can manage the data lifecycle of each storage group from the Actions column on the Data page.

### To manage your data lifecycle:

1. From the Data page, select a Storage Group and click *Actions > Manage Data Lifecycle*.  
The Manage Data Lifecycle dialog loads with the following fields:



The dialog box titled "Manage Data Lifecycle" contains the following fields:

- Data Lifecycle:** A horizontal slider with a red segment and a dotted line.
- Hot Phase:** A label.
- Maximum Age:** A numeric input field with the value "10" and a unit dropdown menu set to "day(s)".
- Rollover Period:** A numeric input field with the value "2" and a unit dropdown menu set to "day(s)".

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

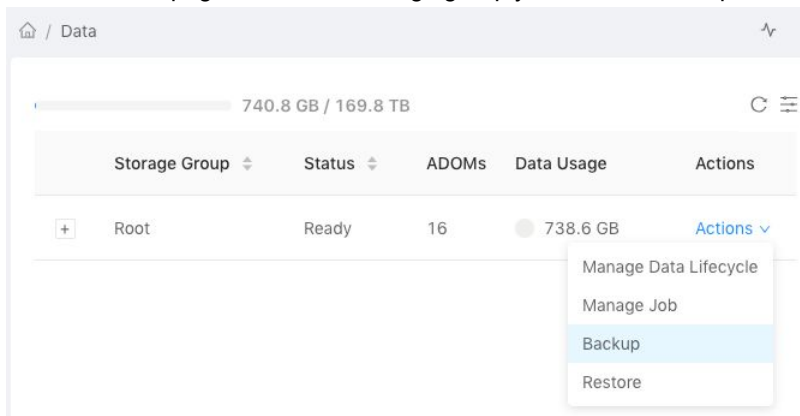
2. In the *Maximum Age* field, select how long you want to store the data in the system. FortiAnalyzer-BigData removes the data from the system after the selected number of days.
3. In the *Rollover Period* field, select the time for log data to roll over to a new partition. FortiAnalyzer-BigData rolls the data into a new partition after the selected number of days.
4. When you are finished, click *OK* to save.

## Data backup



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

### To back up data:

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



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

Field name	Description
<b>HDFS Url</b>	<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 <a href="#">Backup and restore to external HDFS on page 70</a>).</p> </div>
<b>Clean Previous Backup Data</b>	<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>
<b>Backup Timeout</b>	<p>Enter the number of hours before the backup job times out. After the timeout, the job will abort.</p>
<b>Enable Safe Mode</b>	<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 Group 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>
<b>Advanced Config</b>	<p>These configurations define the resources used for the job. Normal users should keep the default configurations.</p>
<b>Enable Scheduled Backup</b>	<p>Enable so the backup can be scheduled automatically.</p>

2. When you are finished, click *Save & Backup* to begin the backup process.
3. You can monitor the status of your backup by navigating to *Jobs > Storage Group 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 group 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 Group 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 Group Backup	
Storage Group :	Root	
HDFS Url :	hdfs://cluster/backup001	
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 Group 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 Group Configuration dialog.

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 group 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 Group 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** ✕

<b>Template :</b>	Storage Group Backup
<b>Storage Group :</b>	Root
<b>HDFS Url :</b>	hdfs://cluster/backup001
<b>Backup Timeout :</b>	24
<b>Clean Previous Backup Data :</b>	false
<b>Enable Safe Mode :</b>	false
<b>Advanced Config :</b>	

3. In the *HDFS Url* field, copy the URL.  
For example: hdfs://cluster/backup/7o7T
4. Go to *Data* and select the same Storage Group 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 Group Configuration dialog.

**Backup Storage Group Configuration** ✕

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

Storage Group: 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 group. Be cautious when selecting your configurations.

---

### To restore data from a backup

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

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

Restore Storage Group Configuration

Storage Group: Root

\* Select Backup:

Storage Backup at Tue Apr 28 23:33:11 GMT 2020

\* HDFS URL:

hdfs://cluster/backup001

Backup Timestamp:

04-28-2020 16:33:11

Restore Timeout:

24

Hour(s)

Enable Safe Mode:

☐

?

Advanced Config:

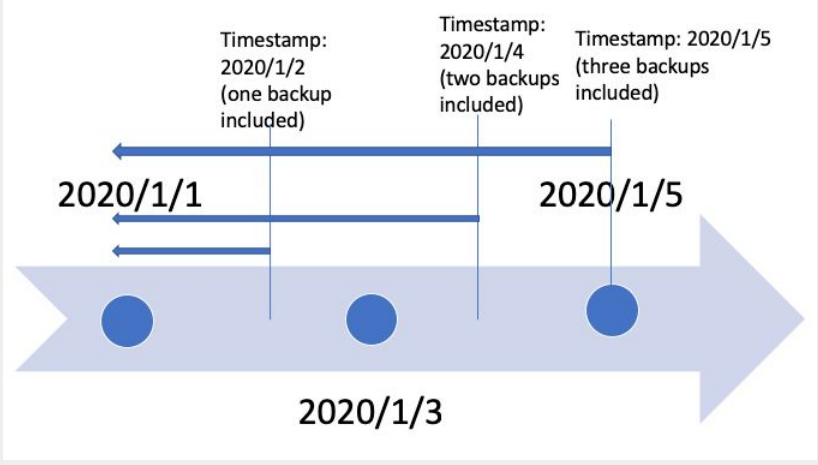

Restore storage group action will pause the log input (Ingestion and Pipeline will be stopped), and **DELETE** all data (**CANNOT** be undone) in current storage group. Please type **restore** to confirm:

Cancel

Restore

- Complete the following fields:

Field name	Description
<b>Select Backup</b>	Select the backup type you used. If the data is from an external system, select Custom.
<b>HDFS URL</b>	Defines the location of the backup. 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.
<b>Backup Timestamp</b>	This config can be used to limit the data that you want to restore. It only applies for multi-backups (multiple incremental backups). The following figure shows how multi-backups are restored:

Field name	Description
	
<b>Enable Safe Mode</b>	<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 Group 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/>
<b>Restore Timeout</b>	Enter the number of hours before the restore job times out. After the hours elapse, the job will abort.
<b>Advanced Config</b>	These configurations define the resources used for the job. Normal users should keep the default configurations.

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 12](#)).
2. Select one of the Security Event Manager host (see [Remotely control blades via CMM on page 18](#)) 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](#)
- 0. [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 -c upgrade` 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 77](#)).



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 77](#)).



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.

```
Please input choice: 8
Please choose method:
1). FTP
0). Cancel
Your choice: 1
Please input server IP [10.106.2.123]:
Please input file path [FAZBD.out]: FAZBD_bootloader.out
Please input username [ftp]:
Please input password:
Your current input:
Server IP: [10.106.2.123]
File path: [FAZBD_bootloader.out]
Username: [ftp]
Password: []
Continue? [y]es/[n]o/[c]ancel: y
```

To upgrade the bootloader of the Security Event Manager Controller, run the following command:

```
fazbdctl -c upgrade -t bootloader
```

To sync all the bootloaders on the Security Event Manager members to the Controller's, run the following command:

```
fazbdctl -c upgrade -t bootloader -h members
```

## 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 `xfs_repair` 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 57](#) and [Data restore on page 62](#).

---



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 -c enable -t 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 Group 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.

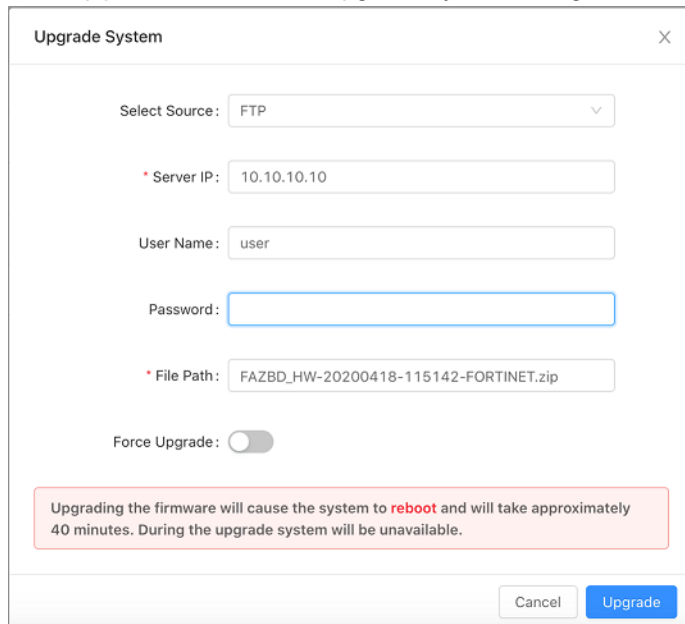
Upgrade takes about 45 minutes. The upgrade starts with the FortiAnalyzer-BigData main host and then the Security Event Manager hosts. During the upgrade, the GUI is not available. Log collecting, LogView, and FortiView operations are also not available.

## To upgrade FortiAnalyzer-BigData via GUI:

1. In the bottom-left of the Navigation bar, click the gear icon .



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




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


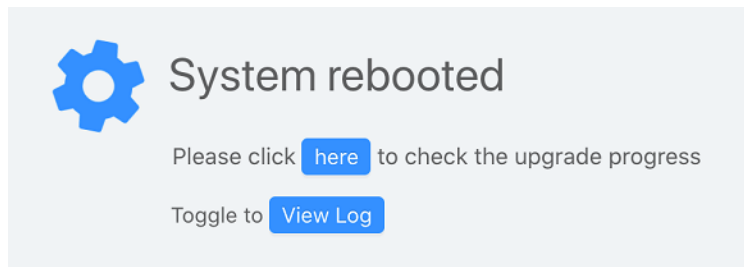


**4.** Click *Upgrade*.


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 79](#).

### To upgrade FortiAnalyzer-BigData via CLI:

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

1. Access the controller blade's Cluster Manger CLI. See [To connect to the Security Event Manager Controller: on page 22](#).
2. Run the following command:  

```
fazbdctl -c upgrade
```
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.
```

# 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. On the extender chassis, power off Blade A1 via the CMM (see [Connect to the Chassis Management Module on page 12](#)).



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 22](#)) and run the following command:  

```
fazbdctl -c set -t appliance -m extender
```
3. Follow the instructions to set a new chassis id for the extender chassis.  
This function updates the chassis id on all hosts, and *hard resets* the FortiAnalyzer-BigData system.
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 -c show -t members
```

  
There should be 12 additional hosts added as members. Wait until all the hosts' status shows as *Joined*.
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.

## 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 22](#)) and run the following command to stop the DHCP service.  
`systemctl stop dhcpd`
2. Run the following command to reset all workers in the Security Event Manager.  
`fazbdctl -c reset -h members -o all-settings`
3. Disconnect the connection between the Internal Switch Modules (Switch Module #1) between the main and extender chassis.
4. Access the Security Event Manager Controller of the main chassis and hard reset FortiAnalyzer-BigData. See [Hard reset FortiAnalyzer-BigData on page 77](#).
5. Access the bootloader of Blade A2 on extender chassis. See [Bootloader on page 65](#).
6. Use the `Set Role` command to set the role to `controller`.
7. Power on Blade A1 of the extender chassis.
8. Access the Security Event Manager Controller of the extender chassis and hard reset FortiAnalyzer-BigData. See [Hard reset FortiAnalyzer-BigData on page 77](#).

# 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 77](#)
- [Hard reset FortiAnalyzer-BigData on page 77](#)

## 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:

1. Access the Security Event Manager Controller (see [To connect to the Security Event Manager Controller: on page 22](#)) and run the following command:

```
fazbdctl -c reset -h 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:

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

1. Access the FortiAnalyzer-BigData Main CLI, and reset the FortiAnalyzer Main host by running the following command:  

```
execute reset [all-except-ip]
```
2. Access the Security Event Manager Controller (see [To connect to the Security Event Manager Controller: on page 22](#)), and run the following command:  

```
fazbdctl -c reset -h cluster -o [all-settings|all-except-ip|all-except-ssh|all-except-ip-ssh]
```

The Security Event Manager Controller reboots after a few minutes.
3. 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 -c show -t members
```
4. After verifying that all the members have a *Joined* status, run the following command to initialize the Security Event Manager:  

```
fazbdctl -c init
```
5. 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"><li>"get image failed"</li><li>"could not find image"</li></ul>	Make sure image from the hosting server is accessible.
An error message displaying: <ul style="list-style-type: none"><li>"checksum verification failed"</li></ul>	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"><li>1. Access the bootloader of the Security Event Manager Controller (see <a href="#">Bootloader on page 65</a>).</li><li>2. Select the "Backup" option to restore the last working OS image to the system.</li></ol>

You can also retry a failed upgrade by using the force flag `-f` in the upgrade command. Enter the following command to forcibly retry upgrading to the same image:

```
fazbdctl -c upgrade -t bd -f
```

## 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"><li>"checksum verification failed"</li><li>"could not find image"</li></ul>	Perform an upgrade with the image of the intended version or latest version.
The Security Event Manager Controller cannot boot up after an upgrade and you cannot connect to the Security Event Manager Controller	Perform the following steps: <ol style="list-style-type: none"><li>1. Access the bootloader of the Security Event Manager Controller (see <a href="#">Bootloader on page 65</a>).</li><li>2. Select the "Backup" option to restore the last working OS image to the system.</li></ol>

Error condition	Troubleshooting suggestion
	<ol style="list-style-type: none"> <li>3. Access the Security Event Manager Controller and perform an upgrade via <code>fazbdctl</code> CLI commands (see <a href="#">Upgrade FortiAnalyzer-BigData on page 72</a>) with the image of the intended version or latest version.</li> <li>4. Rerun the reset command to perform a soft reset.</li> </ol>

## 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 <a href="#">Upgrade FortiAnalyzer-BigData on page 72</a> ) 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"> <li>1. Access the bootloader of the Security Event Manager Controller (see <a href="#">Bootloader on page 65</a>).</li> <li>2. Select the "Backup" option to restore the last working OS image to the system.</li> <li>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.</li> <li>4. Rerun the reset command to perform a hard reset.</li> </ol>

## 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 65](#)).
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 might 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.

There are three types of host roles: Master Node, MetaStore Node and Data Node (see [Roles on page 7](#)). You can find the role type of a host in *Cluster Manager > Hosts*. Some role types require a different method for replacement.

- [Replace a blade with the Data Node role](#)
- [Replace a blade with the Master Node or MetaStore Node role](#)

### To replace a blade with the Data Node role:

1. Access the Security Event Manager Controller (see [To connect to the Security Event Manager Controller: on page 22](#)) and run the following command to decommission the host by its IP address:  
`fazbdctl -c delete -h {member_ip_addr}`
2. Power off the blade, and then remove the blade from the chassis.
3. Insert the replacement blade, and power it on.
4. From the bootloader (see [Bootloader on page 65](#)), set the chassis ID and the blade ID of the replacement blade to match the one from *Cluster Manager > Hosts*.
5. Reconnect to the Security Event Manager Controller and run the following command to ensure that the new blade has joined the cluster:  
`fazbdctl -c show -t members`
6. If the output for the newly added blade shows as "need upgrade", run the following command to upgrade that specific blade:  
`fazbdctl -c upgrade -t bd -h {member_ip_addr}`
7. After the host status changes to "Joined" in the command from step 5, the host will show up in *Cluster Manager > Hosts*.
8. From the Hosts page, click *Assign Role* to add the host.  
The newly added host should have a "new" label.
9. Once the "Assign Role" job finishes running, the blade replacement is done.

### To replace a blade with the Master Node or MetaStore Node role:

1. Go to *Cluster Manager > Services > Actions* and select *Stop All Services*.
2. Power off the blade, and then remove the blade from the chassis.
3. Insert the replacement blade, and power it on.
4. From the bootloader (see [Bootloader on page 65](#)), set the chassis ID and the blade ID of the replacement blade to match *Cluster Manager > Hosts*.
5. Connect to the Security Event Manager Controller (see [To connect to the Security Event Manager Controller: on page 22](#)) and run the following command to ensure that the new blade has joined the cluster:  

```
fazbdctl -c show -t members
```
6. If the output for the newly added blade shows as "need upgrade", run the following command to upgrade that specific blade:  

```
fazbdctl -c upgrade -t bd -h {member_ip_addr}
```
7. After the host status changes to "Joined" in the command from step 5, run the following command to soft reset FortiAnalyzer-BigData (see [Soft reset FortiAnalyzer-BigData on page 77](#)):  

```
fazbdctl -c reset -h cluster
```
8. After the soft reset is finished, the blade replacement is done.

## 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.

There are three types of host roles: Master Node, MetaStore Node and Data Node (see [Roles on page 7](#)). You can find the role type of a host in from *Cluster Manager > Hosts*. Some role types requires a different method for resetting.

- [Reset a blade with a the Data Node role](#)
- [Reset a blade with the Master Node or MetaStore Node role](#)

### To reset a host with Data Node Role:

1. Access the Security Event Manager Controller (see [To connect to the Security Event Manager Controller: on page 22](#)) and run the following command to decommission the host by its IP address:  
`fazbdctl -c delete -h {member_ip_addr}`
2. Access the bootloader of the malfunctioned host (see [Bootloader on page 65](#)), enter the Reset OS option and wait until it finishes rebooting.
3. Reconnect to the Security Event Manager Controller and run the following command to ensure that the host has joined the cluster:  
`fazbdctl -c show -t members`
4. After the host status changes to "Joined", the host will show up in *Cluster Manager > Hosts*.
5. From the Hosts page, click *Assign Role* to add the host.  
The newly added host should have a "new" label.
6. Once the "Assign Role" job finishes running, the host soft reset is done.

### To reset a host with Master Node or MetaStore Node role:

1. Access the bootloader of the malfunctioned host (see [Bootloader on page 65](#)), enter the Reset OS option and wait until it finishes rebooting.
2. Access the Security Event Manager Controller (see [To connect to the Security Event Manager Controller: on page 22](#)) and run the following command to ensure that the host has joined the cluster.  
`fazbdctl -c show -t members`
3. After the host status changes to "Joined", run follow command to soft reset FortiAnalyzer-BigData (see [Soft reset FortiAnalyzer-BigData on page 77](#)):  
`fazbdctl -c reset -h cluster`
4. Once the reset is finishes running, the host soft reset is done.

## 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 33](#)). 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 70](#) 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 section contains 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 Group 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 data lifecycle on page 57](#)). If the data disk on your hosts begin to reach full capacity and are causing the Data Lake services to become unhealthy, you can follow the following steps:


1. From *Cluster Manager > Services*, manually *Stop* the Core service group.
2. Go to *Cluster Manager > Data*, expand the *Root* Storage Group 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. Once you stop receiving messages, go to *Cluster Manager > Services* and manually *Start* the Core service group.

## Change Log

Date	Change Description
2020-08-04	Initial release.
2020-09-21	Replace <i>BigData Cluster</i> with <i>Security Event Manager</i> and general updates.
2020-09-29	Updated <a href="#">Hard reset FortiAnalyzer-BigData on page 77</a> .
2024-01-23	Updated <a href="#">How to scale out on page 75</a> .



**FORTINET®**



Copyright© 2024 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, FortiCare® and FortiGuard®, and certain other marks are registered trademarks of Fortinet, Inc., in the U.S. and other jurisdictions, 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. In no event does Fortinet make any commitment related to future deliverables, features or development, and circumstances may change such that any forward-looking statements herein are not accurate. 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.