



Hewlett Packard
Enterprise

HPE Security ArcSight ADP Event Broker

Software Version: 2.02

Administrator's Guide

July 13, 2017

Legal Notices

Warranty

The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

The network information used in the examples in this document (including IP addresses and hostnames) is for illustration purposes only.

HPE Security ArcSight ArcSight products are highly flexible and function as you configure them. The accessibility, integrity, and confidentiality of your data is your responsibility. Implement a comprehensive security strategy and follow good security practices.

This document is confidential.

Restricted Rights Legend

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notice

© Copyright 2017 Hewlett Packard Enterprise Development, LP

Follow this link to see a complete statement of copyrights and acknowledgements:

<https://www.protect724.hpe.com/docs/DOC-13026>

Support

Contact Information

| | |
|------------------------------|--|
| Phone | A list of phone numbers is available on the HPE Security ArcSight ArcSight Technical Support Page: https://softwaresupport.hpe.com/documents/10180/14684/esp-support-contact-list |
| Support Web Site | https://softwaresupport.hpe.com |
| Protect 724 Community | https://www.protect724.hpe.com |

Contents

| | |
|--|----|
| Chapter 1: Overview | 6 |
| About ADP Event Broker | 6 |
| Event Broker Benefits | 6 |
| Event Broker architecture | 7 |
| Setting up the ADP Event Broker in your environment | 8 |
| Deploying Event Broker | 9 |
| Chapter 2: Producing and Consuming Event Data | 10 |
| Producing Events with SmartConnectors | 10 |
| Pushing JKS files from ArcMC | 11 |
| Consuming Events with ArcSight Investigate & HPE Vertica | 12 |
| Consuming Events with ESM | 12 |
| Consuming Events with Logger | 12 |
| Sending Event Broker Data to Logger | 13 |
| Example Setup with Multiple Loggers in a Pool | 14 |
| Consuming Events with Non-ArcSight Applications | 15 |
| Using Apache Flume to Transfer Events | 15 |
| Consuming Event Broker Events with Apache Hadoop | 16 |
| Architecture for Kafka to Hadoop Data Transfer | 17 |
| Setting Up Flume to Connect with Hadoop | 17 |
| Sample Flume Configuration File | 18 |
| Setting Up Hadoop | 20 |
| Chapter 3: Securing Your Event Broker Deployment | 22 |
| Firewall Configuration | 22 |
| Changing Event Broker security mode | 22 |
| Chapter 4: Managing Event Broker Topics | 24 |
| Default topics | 24 |
| Topic Configuration | 24 |
| Data redundancy and topic replication | 25 |
| Managing topics through ArcMC | 25 |

| | |
|--|----|
| Chapter 5: Managing Event Broker | 26 |
| Managing Event Broker through ArcMC | 26 |
| Enabling Event Broker management through ArcMC | 26 |
| About the Event Broker Manager | 26 |
| Connecting to the Event Broker Manager | 27 |
| Managing Clusters | 28 |
| Viewing Information About a Cluster | 28 |
| Managing Brokers | 29 |
| Viewing Broker Details | 30 |
| Summary | 31 |
| Metrics | 31 |
| Messages count | 31 |
| Per Topic Detail | 31 |
| Managing Topics | 31 |
| Creating Topics | 33 |
| Viewing Topic Details | 34 |
| Topic Summary | 35 |
| Metrics | 36 |
| Operations | 36 |
| Partitions by Broker | 37 |
| Consumers consuming from this topic | 38 |
| Partition Information | 38 |
| Managing Consumers | 38 |
| Viewing Consumer Details | 39 |
| Managing Preferred Replicas | 39 |
| Managing Partitions | 40 |
| Chapter 6: Troubleshooting | 42 |
| Verifying the health of the Event Broker cluster | 42 |
| Diagnosing Common Event Broker Issues | 43 |
| Event Broker Cluster Down | 43 |
| Pod Start Order | 43 |
| Cannot query ZooKeeper | 44 |
| Common Errors/Warnings in ZooKeeper logs | 44 |
| Common Errors/Warnings in Kafka logs | 44 |
| Tuning Event Broker Performance | 46 |
| Increasing Stream Processor EPS | 46 |
| Increasing Kafka retention size/time | 46 |
| Changing the Web Service Admin Password | 47 |

| | |
|---|----|
| Adding a new worker node | 47 |
| Appendix A: The installer.properties file | 48 |
| Glossary | 54 |
| Send Documentation Feedback | 58 |

Chapter 1: Overview

This chapter includes the following topics:

| | |
|---|---|
| • About ADP Event Broker | 6 |
| • Event Broker architecture | 7 |
| • Setting up the ADP Event Broker in your environment | 8 |
| • Deploying Event Broker | 9 |

About ADP Event Broker

The ArcSight Data Platform Event Broker centralizes event processing, enables topic sorting and routing of events, helps you to scale your ArcSight environment, and opens up ArcSight event data to third-party solutions.

It enables you to take advantage of scalable, highly available, multi-broker clusters for publishing and subscribing to event data. The ADP Event Broker integrates with ArcSight Connectors, Logger, and ESM, can be managed and monitored by ArcMC, and is foundational for using ArcSight Investigate.

The ArcSight Data Platform Event Broker provides a packaged version of Confluent Kafka. After you install and configure an Event Broker cluster, you can use ADP SmartConnectors to publish data, and subscribe to that data with ADP Logger, ArcSight ESM, ArcSight Investigate, Apache Hadoop, or your own consumer

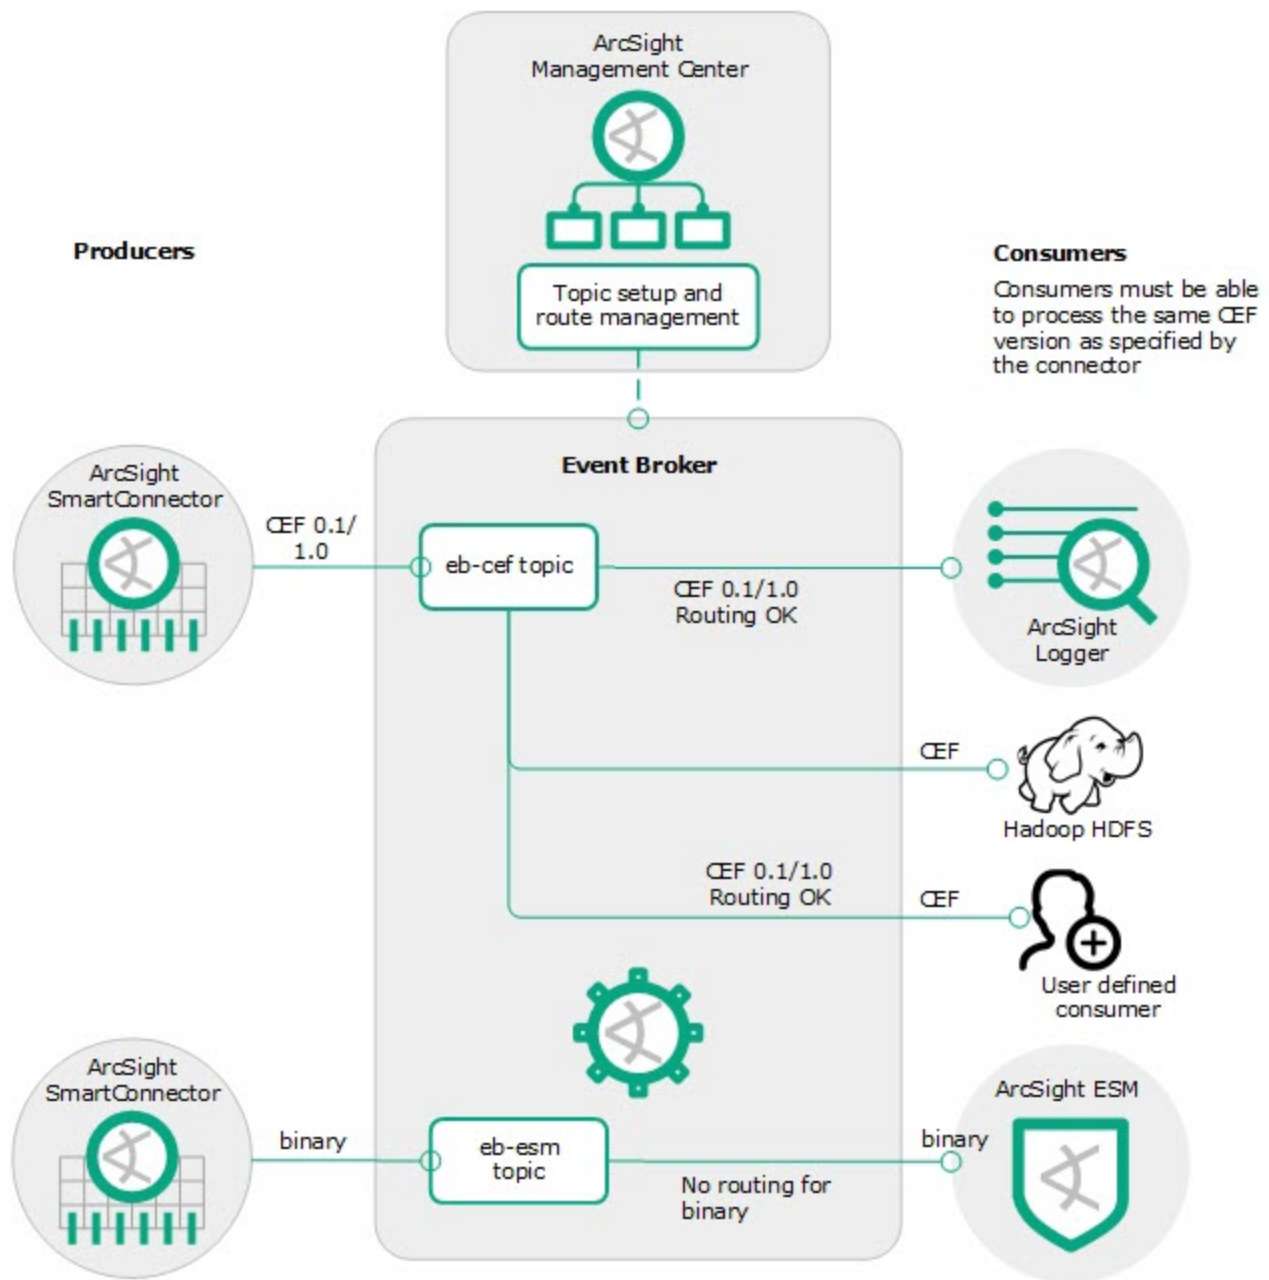
Event Broker Benefits

In addition to open-source Kafka's distributed, high-performance message buses, with resilient redundant message pipelines, Event Broker has many benefits above and beyond open-source Kafka:

- Event Broker is enterprise-ready with these features:
 - **Container-based deployment:** Makes use of Kubernetes deployment for fast central deployment.
 - **Centralized and local management:** Leverage the central management capabilities of ArcSight Management Center, or manage locally with Event Broker Manager.
 - **System and application monitoring:** Monitor Event Broker metrics like your other ArcSight applications, through ArcSight Management Center.
- Event Broker is optimized for deployment with these features:
 - **Ready-to-go security hardening:** Supports TLS, TLS-CA, FIPS, and FIPS-CA.
 - **Event filtering and routing:** Categorize and sort events as needed.

- **Format transformation engine:** Efficiently converts CEF data to Avro format.
- **Ready-to-go topics:** Default topics are installed with the product to get you started quickly.

Event Broker architecture



ArcSight SmartConnectors are the producers that publish data to the ADP Event Broker. Data is sent to Event Broker (in CEF or binary format) to the corresponding topic on the Event Broker.

- You can further route CEF data to specialized topics for categorization.
- Binary data can be used without transformation by ArcSight ESM. Binary topics may not be routed.

Once the data has been processed by Event Broker, you can subscribe to it with producers such as ArcSight Investigate (through Vertica), ArcSight ESM, ADP Logger, Apache HDFS, or your own third-party consumer.

You can manage Event Broker with Event Broker Manager, and with ArcMC, the ArcSight central management and monitoring solution.

Setting up the ADP Event Broker in your environment

Set up and test your Event Broker in a staging environment before deploying to a production environment.

The process of setting up Event Broker in your environment includes the following steps:

1. **Review Event Broker Requirements:** Review the Event Broker technical requirements, and prepare your Kubernetes nodes for Event Broker deployment. Refer to the Event Broker Deployment Guide.
 2. **Deploy Event Broker:** Deploy Event Broker to the designated Kubernetes nodes, as described in the deployment guide.
 3. **Set Up Producers:** Set up one or more SmartConnectors (version 7.6 or later) to produce data for Event Broker. For more information, see the Smart Connectors User Guide.
 4. **Install ArcSight Management Center (ArcMC):** If you plan to manage Event Broker with ArcMC (recommended), install your ArcMC. For more information, refer to the ArcMC Administrator's Guide.
 5. **Configure ArcMC Management:** Configure your Event Broker for management by ArcMC, then add it as a host to ArcMC. For instructions, refer to the ArcMC Administrator's Guide.
 6. **Set Up Consumers:** Set up one or more consumers to consume event data.
- Deploy at least one of the following: ArcSight Investigate, Logger (6.4 or later), ArcSight ESM (6.11.0 or later), Apache Hadoop, or a third-party consumer.
 - Configure consumers to receive events from Event Broker's Kafka cluster.

For more information, see ["Consuming Events with Logger" on page 12](#), ["Consuming Event Broker Events with Apache Hadoop" on page 16](#), or ["Consuming Events with Non-ArcSight Applications" on page 15](#).

Deploying Event Broker

Deploying Event Broker is explained in detail in the ArcSight Installer Deployment Guide, available from the ArcSight documentation repository on [the ArcSight software community](#).

Chapter 2: Producing and Consuming Event Data

Event Broker's publish-subscribe messaging system uses ArcSight SmartConnectors to produce event data, and supports ArcSight Logger and ArcSight ESM, as well as Apache Hadoop and other third-party consumers.

This chapter includes the following topics:

- [Producing Events with SmartConnectors](#)10
- [Consuming Events with ArcSight Investigate & HPE Vertica](#) 12
- [Consuming Events with ESM](#) 12
- [Consuming Events with Logger](#) 12
- [Consuming Events with Non-ArcSight Applications](#)15
- [Consuming Event Broker Events with Apache Hadoop](#)16

Producing Events with SmartConnectors

ArcSight SmartConnectors can publish events to Event Broker Topics. Event Broker supports all SmartConnector types of version 7.6.0 and later.

To publish events you must configure your SmartConnectors to use the Event Broker destination. To send events to multiple topics, you can configure multiple concurrent destinations with the same Event Broker hosts and different topics.

Once configured with an Event Broker destination, the SmartConnector sends events to Event Broker's Kafka cluster, which can then further distribute events to real-time analysis and data warehousing systems. Other applications, including ArcSight Investigate, ESM, Logger, and any third-party application that supports retrieving data from Kafka can receive them, for example, Apache Hadoop.

Tip: You may need to tune your SmartConnectors based on expected throughput.

Event Broker balances events sent by the SmartConnector between nodes by distributing them evenly between the partitions in the configured topic.

Acknowledgments ensure that Event Broker has received the event before the SmartConnector removes it from its local queue. You can disable acknowledgments, require acknowledgment only from the primary replica, or require every replica to acknowledge the event. (Acknowledgments do not indicate that consumers, such as Logger, have received the event data, only that Event Broker itself has.)

The SmartConnector encodes its own IP address as meta-data in the Kafka message for consumers that require that information, such as Logger Device Groups.

For more information about SmartConnectors and how to configure a Event Broker destination, refer to the CEF Destinations chapter of the SmartConnector User's Guide, available for download from the [ArcSight Product Documentation Community on Protect 724](#).

Pushing JKS files from ArcMC

You can push JKS (Java Keystore) files to multiple managed SmartConnectors in ArcMC. First, you will upload the files to a file repository in ArcMC, then push them out to their destination SmartConnectors. You must then configure and enable the Kafka destination on all SmartConnectors.

To upload the Java Keystore files:

1. Prepare the .jks files you want to push and store them in a secure network location.
2. In ArcMC, click **Administration > Repositories > New Repository**.
3. In **Name**, **Display Name**, and **Item Display Name**, enter KAFKA_JKS
4. Enter other required details as needed, then click **Save**.
5. Click **Upload to Repository**.
6. Follow the prompts in the upload wizard and browse to the first .jks file. Note: make sure to choose the individual file option.
7. Upload as many files as needed by repeating the upload wizard.

To push the files to multiple SmartConnectors:

1. In ArcMC , browse to the file repository for the .jks files.
2. Click the **Upload** arrow.
3. Follow the prompts in the wizard and select your destination SmartConnectors.
4. The files are pushed to the managed SmartConnectors and stored in the designated SmartConnector folder.

To configure the Kafka destination on all SmartConnectors:

In ArcMC, click **Node Management > Connectors** tab.

1. Select the SmartConnectors to be configured.
2. Choose **Add a destination** and pick the Kafka destination type.
3. Add the destination details along with the .jks path and password, and save the changes.

Consuming Events with ArcSight Investigate & HPE Vertica

Transformed events in the default topic `eb-internal-avro`, which are in Avro format, can be read by the HPE Vertica database. In turn, once in Vertica storage, event data is accessible for use in ArcSight Investigate searches.

You configure ArcSight Investigate for use with Event Broker as part of the Vertica installer, where you can specify the location of the default Avro topic to which Vertica can subscribe. For instructions, consult the HPE Vertica installer documentation.

Consuming Events with ESM

ArcSight ESM version 6.11.0 or later can subscribe to Event Broker events. ESM requires SmartConnector release 7.6 or later to subscribe to Event Broker topics.

ESM agents are the consumers for Event Broker's publish-subscribe messaging system. An ESM agent can connect to ArcSight Event Broker and consume all events in binary format for the topics it subscribes to.

Additionally, ESM provides data monitors to monitor Event Broker health.

For instruction on configuring ESM 6.11.0 or later as a consumer, see the ESM Administrator's Guide.

Consuming Events with Logger

To subscribe to Event Broker topics with Logger, you must configure an Event Broker receiver on Logger 6.4 or later. Logger's Event Broker receivers are consumers for Event Broker's publish-subscribe messaging system. They receive events in Common Event Format (CEF) from Event Broker's topics. An Event Broker receiver connects to ArcSight Event Broker and consumes all events for the topics it subscribes to.

When configuring an Event Broker receiver, specify the consumer group and topic. You can configure multiple Loggers to consume from the same topic as a part of a consumer group.

For more information about Logger and how to configure an Event Broker receiver, refer to the **Configuration > Receivers** section of the ArcSight Logger Administrator's Guide, available for download from the [ArcSight Product Documentation Community on Protect 724](#).

Sending Event Broker Data to Logger

For a Logger to be able to consume Event Broker events, the Logger must have an Event Broker receiver configured with the Event Broker hosts, consumer group, and event topic list.

SmartConnectors that send data to Event Broker must have an Event Broker destination.

A group of Loggers, called a pool, can be configured to receive and distribute events between themselves. This works similarly to the Logger pool created by using the ArcSight Logger Smart Message Pool destination on SmartConnectors. The difference is that when the SmartConnectors have an ArcSight Logger Smart Message Pool destination, the event load is balanced by each SmartConnector, but when the SmartConnectors have an Event Broker destination, the event load is balanced by the Loggers.

Additional Loggers can be added to the pool simply by configuring the same Event Broker hosts, consumer group, and event topic List in the new Logger's Event Broker receivers, without having to reconfigure either the existing Loggers or any SmartConnectors.

The events retrieved by the Logger pool are distributed among the Loggers in the pool. If one Logger is down, events are rebalanced among existing Loggers. When a Logger is added or removed from the Consumer Group, the event load is distributed across the pool of Loggers.

To send events from a group of SmartConnectors to a pool of Loggers, configure their Event Broker destinations to send events to the topic that the Logger pool is consuming.

To configure Logger to subscribe to event data from specific SmartConnectors, you can do either of the following:

- Configure all the SmartConnectors to publish events to the same topic, then configure the Logger's Event Broker receiver to subscribe to this event topic.
- Configure each SmartConnector to publish events to different topics and then configure the Event Broker receiver on the Logger to subscribe to multiple event topics.

Tip: Loggers in the same Logger pool do not consume the same events, since they are in the same Consumer Group. In high availability situations, you need events to be stored on two different Loggers. To store the same events on two Loggers, configure the Loggers to have different Consumer Group names, but subscribe them to the same event topic.

The number of Loggers in a Logger pool is restricted by the number of event topic partitions. For example, if there are only five partitions, only five Loggers will receive the events. If you have more than five Loggers configured in the same Consumer Group, some Loggers will not normally receive events, but will be available as hot spares. When adding receivers, be sure to increase the number of event topic partitions. See [Managing Topics](#) for more information.

Sending Event Broker data to Logger (Overview):

1. Configure the SmartConnector:

- Setup a SmartConnector to publish to a particular Event Topic. Connectors can only send to a single topic for each destination. Additional destinations need configured if each event needs to go to multiple topics.
Note the number of partitions in the Event Topic.
- Configure the SmartConnector to have an Event Broker destination. For Logger, use an Event Broker destination.
- For more information about SmartConnectors and how to configure an Event Broker destination, refer to the CEF Destinations chapter of the SmartConnector User's Guide, available for download from the [ArcSight Product Documentation Community on Protect 724](#).

2. Configure Logger:

- Create an Event Broker receiver on each Logger in the Logger pool.
- Configure each receiver to subscribe to the Event Topics to which the SmartConnectors are publishing data. To subscribe to multiple topics, indicate the topics by specifying them in the Event Topic List parameter (a list of comma-separated values) while configuring the Event Broker receiver.
- Configure each receiver to be in the same Consumer Group.
For more information on how to configure Logger to receive Kafka events by using Event Broker, see "[Consuming Events with Logger](#)" on page 12, and refer to the Receivers section in the Configuration chapter of the Logger Administrator's Guide, available for download from the [ArcSight Product Documentation Community](#).

Example Setup with Multiple Loggers in a Pool

You can set up your Logger pools to subscribe to events from a particular device type, such as "Firewall." To do this, you would:

1. In ArcMC, create a Kafka topic named *Firewall*.
2. Configure all the SmartConnectors that handle firewall events to publish these events to topic "Firewall."
3. Configure the Loggers in the Logger pool:
 - Create an Event Broker Receiver for each Logger in the pool.
 - Configure the receivers to subscribe to the event topic "Firewall," and include them in the "Logger_Firewall" Consumer Group.

Once all the configuration is set up properly, the Logger pool will subscribe to device type Firewall.

Consuming Events with Non-ArcSight Applications

Event Broker is designed with support for third-party tools. You can create a standard Kafka consumer and configure it to subscribe to Event Broker topics. By doing this you can pull Event Broker events into your own non-ArcSight data lake.

Note: Only use Kafka client libraries that are version 0.10 or later to create consumers.

- All Event Broker nodes, consumers, and producers must be properly configured for DNS/reverse DNS; as well as time, using a time server such as NTP.
- Events are sent in standard CEF (CEF text and CEF binary, which is exclusively for ESM consumption). Anything that can consume from Kafka and understand CEF text can process events.
- You can set up multiple consumer groups, and each group will get a copy of every event. Therefore you can have Logger and Apache Hadoop configured to consume from the same topic and each will get a copy of every event. This enables fanning out multiple copies of events without reconfiguring SmartConnectors or using additional CPU or network resources for them.

Using Apache Flume to Transfer Events

One of the applications you could use to transfer Event Broker events into your data lake is Apache Flume. Flume is designed to push data from many sources to the various storage systems in the Hadoop ecosystem, such as HDFS and HBase. This section describes how to use Apache Flume as a data transfer channel to transfer events from Event Broker to Apache Hadoop or other storage systems.

Prerequisites

- Event Broker installed: Consult the Event Broker Deployment Guide.
- Flume installed: For information on how to install and configure Flume, refer to the Flume documentation, available at <https://flume.apache.org/releases/content/1.6.0/FlumeUserGuide.pdf>.
- Storage system installed: Refer to your storage system documentation.

Procedure

Flume is controlled by an agent configuration file. You must configure Event Broker as the source agent, your storage system as the sink agent, and ZooKeeper as the channel agent in this file.

To configure Event Broker as the source:

Edit the agent configuration file to include the required properties, as in the table below. Configure other properties as needed for your environment.

Required Kafka Source Configuration

| Property | Description |
|------------------|---|
| type | Set to <code>org.apache.flume.source.kafka.KafkaSource</code> . |
| topic | The Event Topic from which this source reads messages. Flume supports only one topic per source. |
| ZooKeeperConnect | The URI of the ZooKeeper server or cluster Kafka is using. POC single node example: <code>zk01.example.com:332181</code> Comma-separated list of nodes in a ZooKeeper cluster example: <code>zk01.example.com:32181,zk02.example.com:32181,zk03.example.com:332181</code> |

To configure the sink:

The required configuration varies. Refer to the Flume documentation for details on your storage system. The section ["Consuming Event Broker Events with Apache Hadoop" below](#) provides an example of how to configure Apache Hadoop as the sink.

Consuming Event Broker Events with Apache Hadoop

Apache Hadoop is a software framework that enables the distributed processing of large data sets across clusters of computers. You can send Event Broker events to Hadoop by using Apache Flume.

This section describes how to set up the Apache Flume agent to transfer Common Event Format (CEF) events from an Event Broker Kafka cluster to Hadoop Distributed File System (HDFS).

It includes the following topics:

- [Architecture for Kafka to Hadoop Data Transfer](#) 17
- [Setting Up Flume to Connect with Hadoop](#) 17
- [Sample Flume Configuration File](#) 18
- [Setting Up Hadoop](#) 20

Architecture for Kafka to Hadoop Data Transfer

Apache Flume uses a source module to read a Kafka topic that has raw CEF events and it then transfers the events using a memory channel, and persist them to HDFS using a sink module. The CEF files are stored on HDFS by time, in a year/month/day/hour directory structure.



Setting Up Flume to Connect with Hadoop

About:

In the simplest deployment model, you need to deploy the Apache Flume agent on a Hadoop node server to pull events, and send them to Hadoop Distributed File System (HDFS).

Prerequisite:

Hadoop must be installed before you can connect it with Flume. If you do not already have your own Hadoop deployment, you can deploy Hadoop on a Red Hat Enterprise Linux 7.2 machine. For more information, see ["Setting Up Hadoop" on page 20](#).

Procedure:

1. Log into your Hadoop server as the user "hadoop".
2. Download Flume from the [Apache download site](#).
3. Uncompress the ".gz" file to your preferred deployment directory.
4. In the configuration file, add your ZooKeeper address and port, Kafka topic, and HDFS address and port.

By default, this configuration persists a CEF file every hour. Alternatively, you could choose to roll using events count or file size. If you have high volume of events, HPE ArcSight recommends

using the event count option instead of time, to avoid running out of memory. For more information, refer to [Flume HDFS sink](#) in the Flume Users' Guide.

5. Execute the following commands to create the Hadoop cefEvents directory:

```
hadoop fs -mkdir /opt
hadoop fs -mkdir /opt/hadoop
hadoop fs -mkdir /opt/hadoop/cefEvents
```

6. Create a configuration file in the Flume conf directory, `bin/flume/conf/`, following the template in "[Sample Flume Configuration File](#)" below. In our example we named the file `kafka.conf`. You can name it whatever is appropriate.

- a. Copy `flume-env.sh.template` as `flume-env.sh`.
- b. Edit `flume-env.sh` file and make the following changes:
 - Set `JAVA_HOME` to the directory where Java was installed on your system.
 - Uncomment the line for `JAVA_OPTS`:

```
export JAVA_OPTS="-Xms100m -Xmx2000m -Dcom.sun.management.jmxremote"
```

- Set `FLUME_CLASSPATH=<Flume install directory>/lib`.

- c. Copy the common jar files from the Hadoop install directory to Flume lib directory:

```
cp <Hadoop install directory>/share/hadoop/common/*.jar /<Flume Install directory>/lib
```

```
cp <Hadoop install directory>/share/hadoop/common/lib/*.jar /<Flume Install directory>/lib
```

- d. Copy `hadoop-hdfs-2.7.2.jar` from Hadoop install directory to Flume lib directory.

```
cp <Hadoop install directory>/share/hadoop/hdfs/hadoop-hdfs-2.7.2.jar /<Flume Install directory>/lib
```

7. Execute the following command to start Flume from its home directory:

```
bin/flume-ng agent --conf conf/ --conf-file conf/kafka.conf --name tier1 -
Dflume.root.logger=INFO,console
```

8. After you start Flume, you can find the files on HDFS by running the following command:

```
hadoop fs -ls -R /opt/hadoop/cefEvents
```

This path has to match the HDFS directory path, created in the Hadoop configuration section.

The files are stored in following structure: "year/month/day/hour".

Sample Flume Configuration File

Before starting Apache Flume, create a configuration file based on the template below.

The configuration file should reside in `bin/flume/conf/`. This file is called `kafka.conf` in our example. You can name your own configuration file whatever is appropriate.

```
#####
```

```
#Sample Flume/Kafka configuration file

#####

#defines Kafka Source, Channel, and Destination aliases
tier1.sources = source1
tier1.channels = channel1
tier1.sinks = sink1

#Kafka source configuration
tier1.sources.source1.type = org.apache.flume.source.kafka.KafkaSource
tier1.sources.source1.ZooKeeperConnect = <ZooKeeperAddress>:Port
# Example Address of ZooKeeper on an Event Broker master node, with port
32181:
# masterNodeIP:32181
tier1.sources.source1.topic = <Kafka_topic>
tier1.sources.source1.groupId = flume
tier1.sources.source1.channels = channel1
tier1.sources.source1.interceptors = i1
tier1.sources.source1.interceptors.i1.type = timestamp
tier1.sources.source1.kafka.consumer.timeout.ms = 150
tier1.sources.source1.kafka.consumer.batchsize = 100
#Kafka Channel configuration
tier1.channels.channel1.type = memory
tier1.channels.channel1.capacity = 10000
tier1.channels.channel1.transactionCapacity = 1000

#Kafka Sink (destination) configuration
tier1.sinks.sink1.type = hdfs
tier1.sinks.sink1.channel = channel1
tier1.sinks.sink1.hdfs.path = hdfs://localhost:9000/opt/\
hadoop/cefEvents/year=%y/month=%m/day=%d
tier1.sinks.sink1.hdfs.rollInterval = 360
```

```
tier1.sinks.sink1.hdfs.rollSize = 0
tier1.sinks.sink1.hdfs.rollCount = 0
tier1.sinks.sink1.hdfs.fileType = DataStream
tier1.sinks.sink1.hdfs.filePrefix = cefEvents
tier1.sinks.sink1.hdfs.fileSuffix = .cef
tier1.sinks.sink1.hdfs.batchSize = 100
tier1.sinks.sink1.hdfs.timeZone = UTC
```

Setting Up Hadoop

This is an overview of the steps necessary to install Apache Hadoop 2.7.2 and set up a one-node cluster. For more information, see <https://hadoop.apache.org/docs/r2.7.2/hadoop-project-dist/hadoop-common/SingleCluster.html>, or refer to the Hadoop documentation for your version.

To install Hadoop:

1. Be sure that your environment meets the Operating System and Java prerequisites.
2. Add a hadoop user.
3. Download and unpack Hadoop.
4. Configure Hadoop for pseudo-distributed operation.
 - Set the environment variables.
 - Set up passphraseless SSH.
 - Optionally, set up Yarn. (You will not need Yarn if you want to use Hadoop only storage and not for processing.)
 - Edit the Hadoop configuration files to set up a core location, a Hadoop Distributed File System (HDFS) location, a replication value, a NameNode and a DataNode.
 - Format the Name node.
5. Start the Hadoop server using the tools provided.
6. Access Hadoop Services in a Browser and login as the user "hadoop".
7. Execute the following commands to create the Hadoop cefEvents directory:

```
hadoop fs -mkdir /opt
hadoop fs -mkdir /opt/hadoop
hadoop fs -mkdir /opt/hadoop/cefEvents
```
8. Execute the following commands to grant permissions for Apache Flume to write to this HDFS

```
hadoop fs -chmod 777 -R /opt/hadoop
hadoop fs -ls
```

9. Execute the following command to check Hadoop system status:
`hadoop dfsadmin -report`
10. Execute the following command to view the files transferred by Flume to Hadoop.
`hadoop fs -ls -R /`

Chapter 3: Securing Your Event Broker Deployment

You are responsible for configuring your Event Broker environment securely according to your business needs and requirements. To help you do this, the Event Broker supports Transport Layer Security (TLS). Ensure that you have your firewalls configured appropriately for your business needs and deployment.

This chapter includes the following topics:

- [Firewall Configuration](#) 22
- [Changing Event Broker security mode](#) 22

Firewall Configuration

You can configure your firewall rules to allow access to only the services that are required.

The Event Broker environment requires the following access:

- Kafka uses port 9093, which is TLS-enabled. All customer data is secured by TLS. (If you are using the Vertica database, you must make port 9092 reachable by all Event Broker nodes, consumers, and producers, but 9092 is not TLS-enabled.)
- ZooKeeper uses port 332181, which must be reachable between all Event Broker nodes.
- The Event Broker Manager uses port 9999 and 10000 to monitor Kafka. These ports must be mutually reachable between all Event Broker nodes.

By default, Kafka ZooKeepers do not use TLS or FIPS to communicate with each other. This communication does not include customer data.

Changing Event Broker security mode

You should decide on a security mode before deployment and setup. In general, the security mode of systems connected to Event Broker (consumers and producers) must be the same as the Event Broker security mode.

TLS is the default configuration. By editing the [installer.properties](#) file, you can enable TLS+CA, as well as FIPS.

You can change Event Broker security modes after deployment, but there will be downtime for Event Broker and its associated systems. You will need to make sure all Event Broker-associated systems are re-configured as well.

Note: Vertica Scheduler does not support TLS. If connecting to Vertica, you should open port 9092, a non-TLS port, on all Event Broker nodes, and only allow access from the Vertica nodes.

To change security mode (Overview):

1. Stop SmartConnectors from sending events. This will close connections. See the SmartConnector User's Guide for information on stopping SmartConnectors from sending events.
2. Stop all consumers (ArcSight Logger, ArcSight ESM, Vertica Scheduler) from consuming from topics in Event Broker. (There is no need to clear out existing messages from the topic.)
3. If the mode change requires that Event Broker consumer or Event Broker producer restarts, then it must disconnect from Event Broker first. See the consumer or producer documentation. Vertica scheduler does not support different security modes.
4. Undeploy the Event Broker containers
5. In a text editor, change the [installer.properties](#) configuration settings.
 - `predeploy.eb.init.client-auth=false`. Set to true to enable TLS+CA.
 - `predeploy.eb.init.fips=false`. Set to true to enable FIPS.

Note: See the [Appendix](#) for a complete list of `installer.properties` settings.

6. Redeploy the Event Broker containers.
7. Follow the consumer and producer documentation to reconfigure those applications to align their security modes as Event Broker.
8. Reconnect the consumers and producers. See the respective product documentation for the steps.

Chapter 4: Managing Event Broker Topics

You can manage your Event Broker topics through Event Broker Manager or through ArcMC.

This chapter includes the following topics:

- [Default topics](#)24
- [Topic Configuration](#)24
- [Data redundancy and topic replication](#)25
- [Managing topics through ArcMC](#)25

Default topics

Event Broker is deployed with several default topics. You can use one of these or configure your own. Topic names are case-sensitive.

| Default Topic Name | Description |
|--------------------------------------|--|
| eb-esm | Supports ESM as a consumer. Use for all ESM events. You can add other topics for ESM events, but cannot create routes for them. |
| eb-cef | Use for CEF 0.1 or CEF 1.0 events. You can create routes to filter events in this topic. |
| eb-other | Use this topic for custom-created consumers. |
| eb-internal-stream-processor-metrics | This topic is for internal use only. Do not configure your SmartConnector destinations to send events to this topic. |
| eb-internal-avro | This topic is for internal use only. Do not configure your SmartConnector destinations to send events to this topic. |
| __consumer_offsets | This topic is for internal use only. Do not configure your SmartConnector destinations to send events to this topic. |
| _schemas | This topic is for internal use only. Do not configure your SmartConnector destinations to send events to this topic. |
| eb-internal-datastore | This topic is for internal use only. Do not configure your SmartConnector destinations to send events to this topic. |

Topic Configuration

HPE ArcSight recommends using different topics for categorization, for example, firewall events in one topic and anti-virus events in another.

- Configure topics based on data isolation requirements or categorization. If you route only to categorized topics, then events are not sent to Vertica for use by ArcSight Investigate.
- Configure partition count for your topics based on throughput and number of consumers. The partition count should be at least equal to the total number of present (and future) consumers.
- Configure the replication factor based on how important the events are. The recommended replication factor for new topics is 2. While you can replicate every topic to every node in the cluster, this is not recommended because the extra traffic reduces throughput and increases disk space requirements.

Data redundancy and topic replication

When setting up Event Broker, you can specify the number of copies (replicas) of each topic Event Broker should distribute.

Kafka automatically distributes each event in a topic to the number of broker nodes indicated by the topic replication level specified during Event Broker configuration. While replication does decrease throughput slightly, HPE ArcSight recommends that you configure a replication factor of at least 2. You need at least one node for each replica. For example, a topic replication level of 5 requires at least five nodes; one replica would be stored on each node.

A topic replication level of 1 means that only one broker will receive that event. If that broker goes down, that event data will be lost. However, a replication level of 2 means that two broker nodes will receive that same event. If one goes down, the event data would still be present on the other, and would be restored to the first broker node when it came back up. Both broker nodes would have to go down simultaneously for the data to be lost. A topic replication level of 3 means that three broker nodes will receive that event. All three would have to go down simultaneously for the event data to be lost.

When you add new consumers, you don't need to update your producers. Event Broker handles the distribution and replication for you.

Refer to the [Apache Kafka documentation](#) for more information.

Managing topics through ArcMC

You can use ArcMC to view and create topics for routing, as well as to create routes, which direct events into appropriate topics.

A *route* is a rule that directs Event Broker to duplicate events that meets certain criteria from a source topic to the route's destination topic. Rules are defined using event field names and expected values.

Using ArcMC, you can view, create, edit and delete routes based on CEF fields and event metadata. (You must create topics before you can route events to them.)

Refer to the ArcMC Administrator's Guide for more information.

Chapter 5: Managing Event Broker

You can manage topic routing and Event Broker infrastructure through ArcMC. Additionally, ArcSight Event Broker provides the Event Broker Manager, a version of Yahoo Kafka Manager, to help you monitor and manage its Kafka services.

For more information about Yahoo Kafka Manager, refer to <https://github.com/yahoo/kafka-manager>.

For more information about Kafka monitoring, refer to the [monitoring section of the Apache Kafka documentation](#).

This chapter includes the following topics:

- [Managing Event Broker through ArcMC](#) 26
- [About the Event Broker Manager](#) 26

Managing Event Broker through ArcMC

You can create topics and routing rules, monitor Event Broker metrics, and receive notifications about Event Broker status through ArcSight Management Center (ArcMC).

Monitored Event Broker parameters include CPU usage, memory, disk usage, throughput, EPS (Events per Second) In, event parsing errors, stream processing EPS, and stream processing lag.

Enabling Event Broker management through ArcMC

To enable Event Broker management in ArcMC, add your Event Broker as a host to ArcMC. The procedure for adding Event Broker as a host is explained in detail in the ArcMC Administrator's Guide, available from [the ArcSight software community](#). The Administrator's Guide also explains in detail how to manage topics, routing rules, monitored metrics, and enabling notifications.

About the Event Broker Manager

The Event Broker Manager enables you to manage your clusters, topics, and partitions. It enables the following monitoring and management options:

- Viewing and managing cluster states, including topics, consumers, offsets, broker nodes, replica distribution, and partition distribution.
- Creating and updating topics.

- Generating partitions and adding partitions to a topic.
- Reassigning partitions to other broker nodes, such as replacing a failed node with a new one.
- Reassigning partition leaders to their preferred broker node after a node temporarily leaves the cluster (for example, in case of a reboot).
- Managing JMX polling for broker-level and topic-level metrics.

Connecting to the Event Broker Manager

Only users that can log into the Event Broker server can access the Event Broker Manager. These users can access the Event Broker Manager by using their local web browser directly from any of the Event Broker nodes or by using SSH forwarding from the Kubernetes worker node where Kafka is running.

You can connect to the Event Broker Manager with most browsers, including Chrome, Firefox and Internet Explorer. For a list of browsers supported in this release, refer to the ADP Support Matrix, available for download from the [ArcSight Product Documentation Community on Protect 724](#).

To access Event Broker Manager:

1. On an Event Broker node, run the command `kubectl get service` to get the list of services.
2. Locate the Event Broker Manager service `eb-kafkamgr-svc` and note its IP and port number.

To connect directly from an Event Broker server node:

1. Log into the Event Broker server.
2. With a supported browser, connect by using the IP and port of the Event Broker manager (as shown previously).

`http://<Event Broker IP:Port>`

Once you connect, the browser displays the Clusters page. See ["Managing Clusters" on the next page](#).

To connect from your local machine:

1. From your local system, set up SSH forwarding and connect by using a command like the following:
`ssh -L <Event Broker port>:<Event Broker IP:port> eb1.example.com`
2. With a supported browser, connect by using the following URL:

`http://<127.0.0.1:Port>`

Once you connect, the browser displays the Clusters page. See ["Managing Clusters" on the next page](#).

Managing Clusters

The **Clusters** page is the Event Broker Manager's home page. From here you can modify, disable or delete a cluster from view in the Event Broker Manager (the cluster itself is not deleted), or drill down into the cluster for more information.

Location: Clusters

Click the *Cluster Name* link. The Event Broker Manager displays the Cluster Summary page. See "[Viewing Information About a Cluster](#)" below.

To edit the cluster:

1. Click **Modify**. The Event Broker Manager displays the **Update Cluster** page.
2. Update the appropriate fields, and click **Save**.

Editing the cluster is an advanced operation, and normally the cluster should never be edited.

To disable the cluster:

Click **Disable**. Once a cluster has been disabled, a **Delete** button is displayed.

To delete the cluster:

Click **Delete**.

Viewing Information About a Cluster

On the **Summary** page, you can view the ZooKeeper processes in your cluster and drill down into its topics and broker nodes for more information.

Location: Clusters > *Cluster Name* > Summary

Event Broker **event-broker** Cluster ▾ Brokers Topic ▾ Preferred Replica Election Reassign Partitions Consumers

Clusters / event-broker / Summary

Cluster Information

| | |
|------------|--|
| Zookeepers | 192.168.1.100:2181, 192.168.1.101:2181, 192.168.1.102:2181 |
| Version | 0.10.0.0 |

Cluster Summary

| | | | |
|--------|---|---------|---|
| Topics | 8 | Brokers | 3 |
|--------|---|---------|---|

Topics Link Brokers Link

To view information about your cluster:

- If the cluster is not yet open, click **Cluster > List** in the navigation bar. Then click the *Cluster Name* link.
- If the cluster is already open, click **Clusters > Cluster Name > Summary**

To view or edit the topics in your cluster:

Click the **Topics** link. See "[Managing Topics](#)" on page 31.

To view or edit the broker nodes in your cluster:

Click the **Brokers** link. See "[Managing Brokers](#)" below.

Managing Brokers

On the **Brokers** page, you can view overview information on all of your broker nodes and drill down into a broker for more information.

Note: The term *Brokers* is used synonymously with *Event Broker Nodes* in ArcMC. Both terms describe a single node running Kafka.

Location: Clusters > *Cluster Name* > Brokers

To view the broker nodes in your cluster:

Click **Brokers** in the navigation bar. The **Brokers** page opens.

To see more information about a specific broker:

Click the broker's *Id* link. The *Broker Name* ID opens. See "[Viewing Broker Details](#)" below.

Viewing Broker Details

You can view detailed information about a broker from the *Broker Name* details page.

Location: Clusters > *Cluster Name* > Brokers > *Broker Name*

To view information on a specific broker:

1. Click **Brokers** in the navigation bar.
2. Click the *Broker Name* link. The *Topic Name* page opens.

From here you can view the following:

- "[Summary](#)" on the next page
- "[Metrics](#)" on the next page
- "[Messages count](#)" on the next page
- "[Per Topic Detail](#)" on the next page

Summary

In the **Summary** section, you can see an overview of your broker, including the number of topics and partitions located on it.

Metrics

In the **Metrics** section, you can view information about the data flow.

Messages count

In the **Messages** section, you can view a message view chart.

Per Topic Detail

In the **Per Topic Detail** section, you can view topic replication and partition information and drill down to view more information on each topic.

To see more information about a specific topic:

Click the *Topic Name* link in the **Per Topic Details** section. See ["Viewing Topic Details" on page 34](#).

Managing Topics

On the **Topics** page, you can run or generate partition assignments, add a new partition, and drill down into individual topics for more information.

Location: Clusters > *Cluster Name* Topic > List

Event Broker **event-broker** Cluster Brokers Topic Preferred Replica Election Reassign Partitions Consumers

Clusters / event-broker / Topics

Operations

Generate Partition Assignments Run Partition Assignments Add Partitions

Topics

Show 10 entries Search:

| Topic | # Partitions | # Brokers | Brokers Spread % | Brokers Skew % | # Replicas | Under Replicated % | Producer Message/Sec |
|------------------------------------|--------------|-----------|------------------|----------------|------------|--------------------|----------------------|
| 7864-connector | 10 | 3 | 100 | 0 | 3 | 0 | 0.00 |
| __consumer_offsets | 50 | 3 | 100 | 0 | 2 | 0 | 0.00 |
| firewall | 10 | 3 | 100 | 0 | 3 | 0 | 0.00 |
| syslog-topic | 30 | 3 | 100 | 0 | 2 | 0 | 0.00 |

Note: These topics are installed by default:

- `__consumer_offsets`
- `_schemas`
- `eb-internal-datastore`
- `eb-internal-stream-processor-metrics`.

These are used internally by Event Broker and should not be modified.

To manage the topics in your cluster:

Click **Topic > List** in the navigation bar.

To view information on a topic:

Click the *Topic Name* link. The *Topic Name* page displays the topic's summary, metrics, consumers, and partitions. See "[Viewing Topic Details](#)" on [page 34](#).

To generate partition assignments:

1. Click **Generate Partition Assignments**.
2. Select the topics and broker nodes to reassign.
3. Click **Generate Partition Assignments**.

To assign partitions as generated:

1. Click **Run Partition Assignments**.
2. Select the topics to reassign.
3. Click **Run Partition Assignments**.

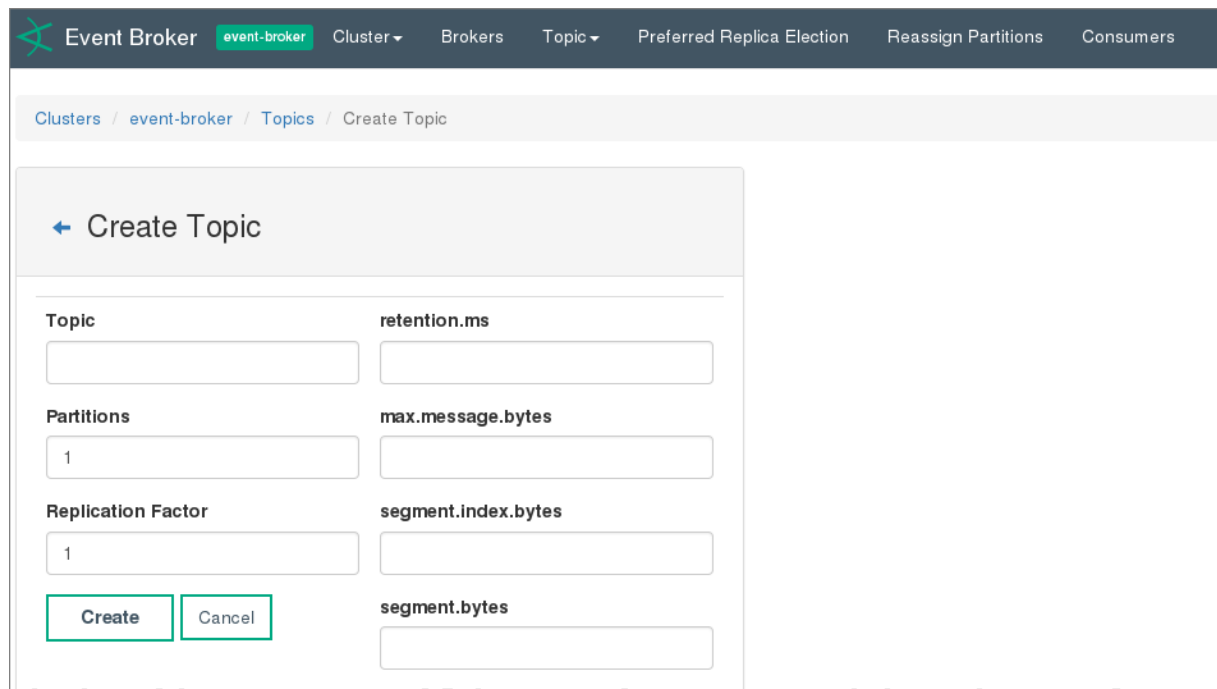
To add a partition:

1. Click **Add Partition**.
2. Enter the new number of partitions.
3. Select the topics and broker nodes.
4. Click **Add Partitions**.

Creating Topics

You can create a new topic on the **Create Topic** page.

Location: Clusters > *Cluster Name* Topics > Create Topic



The screenshot shows the 'Create Topic' page in the Event Broker interface. The navigation bar at the top includes 'Event Broker', 'event-broker', 'Cluster', 'Brokers', 'Topic', 'Preferred Replica Election', 'Reassign Partitions', and 'Consumers'. Below the navigation bar, a breadcrumb trail reads 'Clusters / event-broker / Topics / Create Topic'. The main content area is titled 'Create Topic' and contains a form with the following fields:

| | |
|--------------------|---------------------|
| Topic | retention.ms |
| Partitions | max.message.bytes |
| Replication Factor | segment.index.bytes |
| | segment.bytes |

At the bottom left of the form, there are two buttons: 'Create' and 'Cancel'.

Note: You cannot delete topics once they have been created.

To open the Add Topic page:

Click **Topic > Create** in the navigation bar.

To create a new topic:

Fill in the fields and click **Create**.

Viewing Topic Details

You can see details about a topic, including information about the summary, metrics, consumers, and partitions from the *Topic Name* details page.

Location: Clusters > *Cluster Name* Topics > *Topic Name*

To view information on a specific topic:

1. Click **Topic > List** in the navigation bar.
2. Click the *Topic Name* link. The *Topic Name* page opens.

From here you can view the following:

- ["Topic Summary" on the next page](#)
- ["Metrics" on page 36](#)
- ["Operations" on page 36](#)
- ["Partitions by Broker" on page 37](#)
- ["Consumers consuming from this topic" on page 38](#)
- ["Partition Information" on page 38](#)

Topic Summary

In the **Topic Summary** section, you view information on the topic's replicas, partitions, and broker nodes.

| Topic Summary | |
|-----------------------------|--------|
| Replication | 3 |
| Number of Partitions | 10 |
| Sum of partition offsets | 0 |
| Total number of Brokers | 3 |
| Number of Brokers for Topic | 3 |
| Preferred Replicas % | 100 |
| Brokers Skewed % | 0 |
| Brokers Spread % | 100 |
| Under-replicated % | 0 |
| Config | Value |
| cleanup.policy | delete |

Metrics

In the **Metrics** section, you can view information about the data flow.

| Metrics | | | | |
|-----------------------------|-------|-------|-------|--------|
| Rate | Mean | 1 min | 5 min | 15 min |
| Messages in /sec | 23.14 | 28.80 | 28.80 | 28.80 |
| Bytes in /sec | 2.1k | 2.6k | 2.6k | 2.6k |
| Bytes out /sec | 6.4k | 10k | 10k | 10k |
| Bytes rejected /sec | 0.00 | 0.00 | 0.00 | 0.00 |
| Failed fetch request /sec | 0.00 | 0.00 | 0.00 | 0.00 |
| Failed produce request /sec | 0.00 | 0.00 | 0.00 | 0.00 |

Operations

In the **Operations** section, you can reassign partitions, generate partition assignments, add partitions, update the topic configuration, and manually assign topics to broker nodes.

| Operations | | |
|---------------------|--------------------------------|------------------------------|
| Reassign Partitions | Generate Partition Assignments | |
| Add Partitions | Update Config | Manual Partition Assignments |

To reassign partitions:

Click **Reassign Partitions**.

To generate partition assignments:

1. Click **Generate Partition Assignments**.
2. Select the topics and broker nodes to reassign.
3. Click **Generate Partition Assignments**.

To add a partition:

1. Click **Add Partitions**.
2. Enter the new number of partitions.
3. Select the topics and broker nodes.
4. Click **Add Partitions**.

To update the topic's configuration:


1. Click **Update Config**.
2. Edit the configuration fields.
3. Click **Update Config**.

To specify partition assignments:

1. Click **Manual Partition Assignment**.
2. Select the desired assignments.
3. Click **Save Partition Assignment**.

Partitions by Broker

In the **Partitions by Broker** section, you can see topic partition information and drill down to see details for each broker.

| Partitions by Broker | | | |
|---|-----------------|-----------------------|---------|
| Broker | # of Partitions | Partitions | Skewed? |
| 1  Broker Link | 10 | (0,1,2,3,4,5,6,7,8,9) | false |
| 2 | 10 | (0,1,2,3,4,5,6,7,8,9) | false |
| 3 | 10 | (0,1,2,3,4,5,6,7,8,9) | false |

To view details on a broker:

Click the Broker link. The Topic Summary page displays information on the topic's lag, partitions, and consumer offset.

In Kafka Manager, users will see different offset values between Binary (ESM) and CEF (Investigate or Logger) topics. In CEF topics, the offset value can generally be associated with number of events that passed through the topic. Each message is an individual event. However, that same association cannot be made in Binary topics.

Consumers consuming from this topic

In the **Consumers consuming from this topic** section, you can drill down to see details on each consumer.

New consumers can take some time to display properly. Give the process time to populate correct data.

To view details on a consumer:

Click the *Topic Name* link. The Topic Summary page displays information on the topic's lag, partitions, and consumer offset.

Partition Information

In the **Partition Information** section, you can view information about the topic's partitions and drill down for more information on each leader.

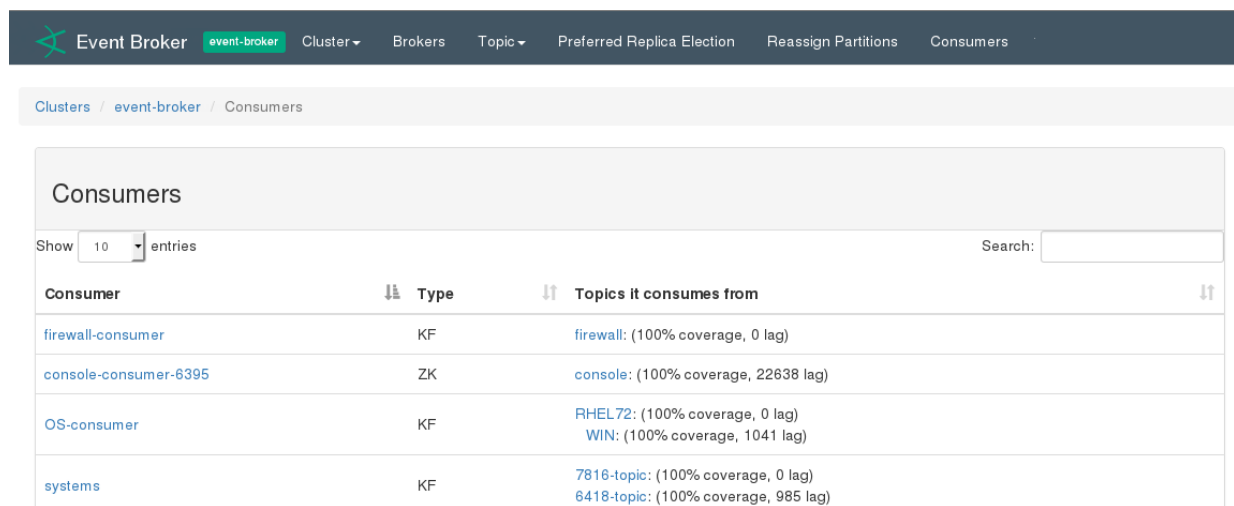
To view details on a Leader:

Click Leader link. The *Broker Name* ID page displays the broker's summary, metrics, message count, and topic details. See ["Viewing Broker Details" on page 30](#).

Managing Consumers

On the **Consumers** page, you can see a list of consumers, view their type, the topics they consume, and drill down into each consumer and topic for more information.

Location: Clusters > *Cluster Name* > Consumers



| Consumer | Type | Topics it consumes from |
|-----------------------|------|--|
| firewall-consumer | KF | firewall: (100% coverage, 0 lag) |
| console-consumer-6395 | ZK | console: (100% coverage, 22638 lag) |
| OS-consumer | KF | RHEL72: (100% coverage, 0 lag) WIN: (100% coverage, 1041 lag) |
| systems | KF | 7816-topic: (100% coverage, 0 lag) 6418-topic: (100% coverage, 985 lag) |

To view or edit the consumers in your cluster:

Click **Consumers** in the navigation bar.

To view more details on a specific consumer:

Click the *Consumer Name* link. The *Consumer Name* page displays details about the consumer. You can drill down further for more information.

To view more details on the topic it consumes:

Click the *Topic Name* link. The *Topic Name* page displays details about the topic. You can drill down further for more information.

Viewing Consumer Details

You can see a information about a consumer and drill down on the topics it consumes from the *Consumer Name* details page.

Location: Clusters > *Cluster Name* Consumer > *Consumer Name*

To view information on a consumer:

1. Click Clusters > *Cluster Name* Consumer.
2. Click the *Consumer Name*.

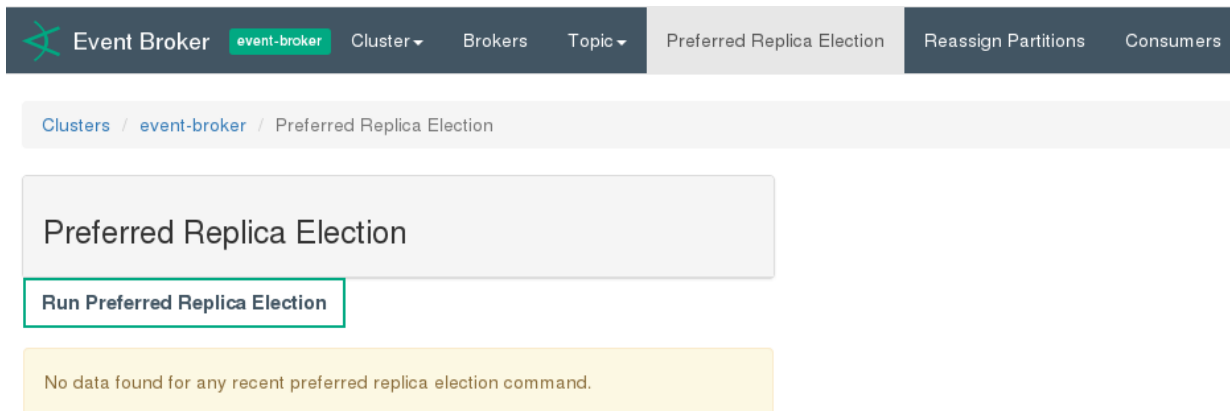
To view information on the consumed topic:

1. Click the *Topic Name*. The Consumed Topic Information page displays information about the topic. Click the topic name for more information.

Managing Preferred Replicas

You can update the replicas for each cluster on the **Preferred Replica Election** page.

Location: Clusters > *Cluster Name* > Preferred Replica Election



The screenshot shows the 'Preferred Replica Election' page. At the top is a navigation bar with 'Event Broker' and a dropdown menu containing 'event-broker', 'Cluster', 'Brokers', 'Topic', 'Preferred Replica Election' (which is highlighted), 'Reassign Partitions', and 'Consumers'. Below the navigation bar is a breadcrumb trail: 'Clusters / event-broker / Preferred Replica Election'. The main content area has a header 'Preferred Replica Election' and a button 'Run Preferred Replica Election'. Below this is a yellow message box stating: 'No data found for any recent preferred replica election command.'

To open the Preferred Replica Election page:

Click **Preferred Replica Election** in the navigation bar.

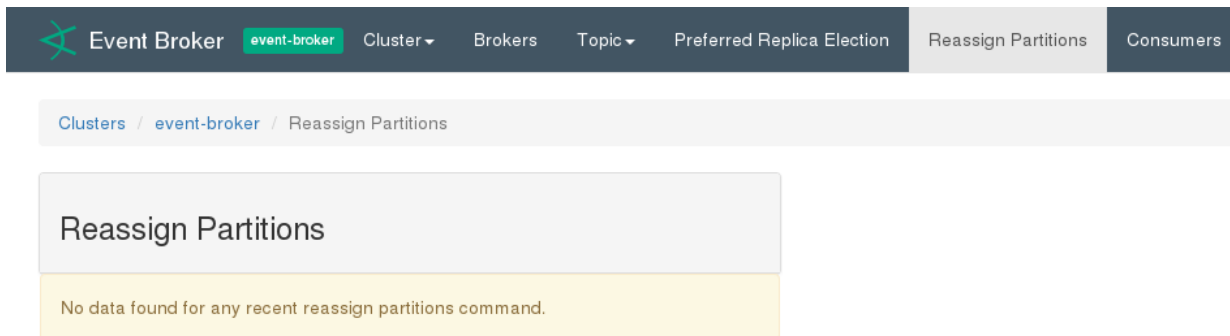
To run the Preferred Replica Election for your topic:

Click **Run Preferred Replica Election**.

Managing Partitions

You can reassign partitions for your cluster on the **Reassign Partitions** page.

Location: Clusters > *Cluster Name* > Reassign Partitions



The screenshot shows the 'Reassign Partitions' page. At the top is a navigation bar with 'Event Broker' and a dropdown menu containing 'event-broker', 'Cluster', 'Brokers', 'Topic', 'Preferred Replica Election', 'Reassign Partitions' (which is highlighted), and 'Consumers'. Below the navigation bar is a breadcrumb trail: 'Clusters / event-broker / Reassign Partitions'. The main content area has a header 'Reassign Partitions' and a yellow message box stating: 'No data found for any recent reassign partitions command.'

To open the Reassign Partitions page:

Click **Reassign Partitions** in the navigation bar.

To reassign the partitions for your topic:

Click **Reassign Partitions**.

Chapter 6: Troubleshooting

These troubleshooting tips may prove helpful in resolving issues with Event Broker.

This chapter includes the following topics:

- [Verifying the health of the Event Broker cluster](#)42
- [Diagnosing Common Event Broker Issues](#)43
- [Tuning Event Broker Performance](#)46

Verifying the health of the Event Broker cluster

Verify the health of each container: run `kubectl get pods -o wide` to list pods and their status.

View Kubernetes logs for each container: run `kubectl logs`

```
# kubectl logs [POD ID/NAME]
```

```
# kubectl logs [WEB SERVICE POD ID/NAME] -c atlas-web-service
```

Verify data flows through the system: check any of the following.

- In ArcMC, review the EPS graph. This indicates whether events are flowing through the stream processor (routing and transforming).
- In Vertica server, check the Kafka scheduler status to see event count and reject count. You should be able to see the event count increasing.

```
# ./install-vertica/kafka_scheduler status
```

- Check the Kafka manager offset with the `select count(*)` in Vertica. The count should increase over time. (for example, `SELECT COUNT (*) FROM investigation.events;`)
- All topics: check the offset for each topic in Event Broker Manager. you should see the value increasing.

Verify that Web Service APIs are healthy:

- Check logs of the web service container (see command above)
- Make sure the port is bound:

```
# netstat -lntp | grep 38080
```

- Verify Vertica Scheduler is running.
- Check the Kafka Scheduler status.

```
# watch ./root/install-vertica/kafka_scheduler status
```

- Check whether the offset is increasing in the status output. If not, then there may no data in the Avro topic, or if Avro contains data there may be a problem.
- Verify the topic partition count and distribution.
- Check that the configured partition count matches its expected value
- Check the partition count or replication factor for the topic using Event Broker Manager.

Diagnosing Common Event Broker Issues

The following can help to diagnose common Event Broker issues.

Event Broker Cluster Down

The number of nodes required to keep an Event Broker cluster operating depends on the replication factor. If the replication factor is only 1, which is not recommended, then all Kafka nodes in the EB cluster need to be up to make the EB cluster function correctly. In general, if the replication factor is N, then the system will tolerate up to N-1 server failures without losing any records committed to the log.

Pod Start Order

After deployment, pods are configured to start in the following order (downstream pods will not start until the dependencies are met.)

1. A quorum of ZooKeeper pods in the cluster must be up (2 of 3, or 3 of 5). The total number of ZooKeepers must be odd.
2. All Kafka pods must be up
3. Schema Registry pod must be up
4. Bootstrap Web Service, Event Broker Manager
5. Transformation Stream Processor, Routing Stream Processor

Cannot query ZooKeeper

This can occur when running the `kubectl get pods` command to get status of the pods, downstream pods (as defined in the pod start order) do not stay up, and status is a 'CrashLoop'-type error.

- Check that ZooKeeper pods are running.
- If the ZooKeeper pod status is Pending, you may not have labeled the nodes (zk=yes). Verify that the nodes are labeled using the `kubectl get nodes -L=zk` command.
- Verify that you configured an odd number of ZooKeepers in the [installer.properties](#) `predeploy.eb.zookeeper.count` attribute.
- Check the ZooKeeper pod logs for errors using the `kubectl logs <pod name>`.

Common Errors/Warnings in ZooKeeper logs

- **Quorum Exceptions:** A leader cannot be elected. If you see this type of error, check the conditions above.
- **Socket error:** this can occur if there are too many connections. The solution is to restart the pod using the `kubectl delete <pod_name>`. The pod will be recreated automatically.

Common Errors/Warnings in Kafka logs

Cannot Register ID: In some cases, a broker node cannot register its ID. This can be caused by multiple broker nodes with the same ID. This is a rare situation that can occur when you are adding and removing nodes from the cluster and you do not define the cluster properly. Connect to each system running a Kafka broker and check the assigned `broker.id` value of each, in `/opt/arcsight/k8s-hostpath-volume/eb/kafka/meta.properties`. The `broker.id` value defined on each Kafka node must be unique.

SSL Connection Errors: These are warnings that occur if there is a connection issue between Kafka and consumer or producer.

Cannot communicate with other brokers: Host names may not be configured properly. It is possible that the node cannot perform reverse lookup or that DNS is not set up properly.

Event Broker default topics not created on first deployment: In this instance, the Bootstrap Web Service log contains 500 response code (the response from the Schema Registry), and topics are not created. Try undeploying the Event Broker containers, and then redeploy them.

One or more connectors cannot send data to Kafka: Check the following:

- The connection configuration is set properly in the connector.
- The encryption mode (TLS, TLS+FIPS, TLS+CA, TLS+FIPS+CA) is the same for both the Connector and Event Broker.
- Make sure you can connect to the Kafka port on the system and that there are no network issues.

Cannot retrieve the certificate error when connecting: Make sure that time is synced across all systems in the data pipeline.

- Check whether the Kafka pod is down. Did you configure the connector with only one broker address and that broker is down; If you expect that there are multiple brokers, they must be all configured in connector as a comma-separated list;
- If the replication factor is set to 1 and a Kafka broker is down, data will not be able to sent through Event Broker. Fix the broker issue to bring it back up. In general, topics should be configured with replication factor greater than 1 so as to prevent this scenario.

Kafka is resyncing: This may cause event throughput slowdown, but will not stop event flow.

Vertica cannot read events from Kafka: After verifying that the Event Broker is still up, check the following:

- For a new setup: Check that Kafka scheduler is configured to communicate to Kafka port 9092. Also, check the network connection.
- For an existing setup (with Vertica consumers): Offset may not be recognized: In this scenario, the Kafka scheduler fails to recognize offset IDs of messages that are in the topic. It can happen if the Kafka scheduler unexpectedly stops reading from the topic, and then is restarted.

Solution: Execute the `kafka_scheduler delete` command to delete the metadata. After doing this, immediately run the `kafka_scheduler create` command to set up the scheduler.

- Existing set up: You have configured all brokers that contain the topic the consumer connects to, and the brokers which are configured for that consumer are down.

An EB component crashes: Check the following:

- Check the container start up order (above). Have any of the dependency pods not started or crashed?
- It could be that the JVMs require more memory that the system has available.
- Check the number of open sockets.

Event Broker EPS is lower than expected: Check resource constraints on Event Broker nodes, such as CPU, memory, or disk space. Also, check usage with ArcMC.

Network bottleneck: In this case, the Stream Processor is not able to keep up with transformation, or is resource-constrained in some way. In ArcMC, the Stream Processor metric will be lower than the connector EPS. Check that you have sufficient resources, memory, CPU.

Continuous network failures: This may be related to the management of TCP/IP resources. TIME_WAIT is the parameter which indicates the amount of time the node will try take to finish closing a connection and the amount of time before it will kill a stale connection. Try reducing the value from its default. Edit the file `/etc/sysctl.conf` and add these lines to the end of it (or edit the existing values):

Decrease TIME_WAIT seconds

```
net.ipv4.tcp_fin_timeout = 10
```

Recycle and Reuse TIME_WAIT sockets more quickly

```
net.ipv4.tcp_tw_recycle = 1
```

```
net.ipv4.tcp_tw_reuse = 1
```

After editing the file you should run

```
$ sysctl --system
```

Tuning Event Broker Performance

The following can help improve the performance of Event Broker.

Increasing Stream Processor EPS

You can increase Stream Processor EPS by adding more stream processor instances using the ArcSight installer configuration UI. The configuration in the ArcSight installer configuration UI affects transforming stream processor only (c2av). It does not change the routing stream processor. You cannot modify the number of streams in the routing stream processor.

When you change this value, you do not need to redeploy Event Broker. Please note that this change will increase the number of pods. You will see this difference when you run the `kubectl get pods` command.

Increasing Kafka retention size/time

You can change the value of retention size or time in any topic using Event Broker Manager after deploying Event Broker containers, and it will be applied immediately. You can change this while events are flowing through the topic.

To change the default values *before* you deploy, change the values in the [installer.properties](#) file.

Changing the Web Service Admin Password

Before you deploy, the web service admin password is changed in the [installer.properties file](#).

Adding a new worker node

To add a new worker node, label the new node (delete or overwrite existing label with a different label). Remove the label from the old node. Kubernetes should start Kafka on the new node. Then, reassign partitions on the new node. Data copying will take some time to complete.

Appendix A: The installer.properties file

The installer.properties file controls several important settings for your Event Broker installation. Those settings are detailed here.

To edit the installer.properties file: open the file in a text editor and make changes as needed.

In order for changed settings to take effect, you will need to undeploy Event Broker and then re-deploy.

| Setting | Notes |
|---|--|
| ## All Event Broker components will use FIPS-certified encryption algorithms | |
| predeploy.eb.init.fips=false | Turns FIPS on. Not recommended to change after deployment. Should undeploy, then redeploy. |
| | |
| ## Event Broker kafka will use TLS Client Authentication to verify client connections | |
| predeploy.eb.init.client-auth=false | Turns TLS-CA on. Not recommended to change after deployment. Should undeploy, then redeploy. |
| | |
| ## Number of partitions for Event Broker default topics in kafka | |

| Setting | Notes |
|--|--|
| <code>predeploy.eb.init.noOfTopicPartitions=5</code> | Default value. Will only affect newly created topics. (Add new partitions to existing topics with the Event Broker Manager.) |
| ## Replication factor for Event Broker default topics in kafka | |
| <code>predeploy.eb.init.topicReplicationFactor=2</code> | Default value. Will only affect newly created topics. (Must delete old topics to change replication factor.) |
| ## kafka log retention size | |
| <code>predeploy.eb.init.kafkaRetentionBytes=10737418240</code> | Default value per partition per topic. Very small, will definitely require customer adjustment. Requires calculation on customer behalf. Deletion will occur when EB hits either the duration or retention bytes, whichever comes first. |

| Setting | Notes |
|--|--|
| ## kafka log retention size for the Vertica Avro topic. This is uncompressed and requires more space to hold events for the same duration. | |
| predeploy.eb.init.kafkaRetentionBytesForVertica=10737418240 | Default value per partition per topic. Very small, will definitely require customer adjustment. Requires calculation on customer behalf. May require additional space than other topics because data is uncompressed. To ensure data retention is the same as other topics, this topic may need to be significantly larger than other topics, as large as a factor of 7 or more. Deletion will occur when EB hits either the duration or retention bytes, whichever comes first. |
| ## kafka log retention duration | |

| Setting | Notes |
|---|--|
| <code>predeploy.eb.init.kafkaRetentionHours=672</code> | Based on environment. Requires calculation on customer behalf. Applies to all topics, including those created through ArcMC. Deletion will occur when EB hits either the duration or retention bytes, whichever comes first. |
| | |
| ## kafka inter-broker protocol version | |
| <code>predeploy.inter.broker.protocol.version=0.10.1.0</code> | Only to be used during upgrades. |
| | |
| ## The message format version the broker will use to append messages to the logs. | |
| <code>predeploy.log.message.format.version=0.10.1.0</code> | Only to be used during upgrades. |
| | |
| ## Size of kafka and ZooKeeper pet-sets | |
| <code>predeploy.eb.kafka.count=3</code> | Determines cluster size for Kafka. Must match number of worker nodes labelled as kafka=yes in K8s. 1 node to 1 host. |

| Setting | Notes |
|---|--|
| <code>predeploy.eb.zookeeper.count=3</code> | Determines cluster size. Max of 7. Must match number of worker nodes labelled as zk=yes in K8s. MUST be an odd number. |
| <code>## Host path to store data persistently</code> | |
| <code>predeploy.eb.kafka.path=/opt/arcsight/k8s-hostpath-volume/eb/kafka</code> | Go big. See sizing guidelines. Will be created if it does not exist. |
| <code>predeploy.eb.zookeeper.path=/opt/arcsight/k8s-hostpath-volume/eb/zookeeper</code> | Will be created if it does not exist. |
| <code>## ArcMC hostname</code> | |
| <code>predeploy.eb.arcmc.hosts=localhost:443</code> | Please ignore. |
| <code>## The endpoint identification algorithm to validate the server hostname using the server certificate.</code> | |
| <code>predeploy.ssl.endpoint.identification.algorithm=https</code> | If reverse DNS is not set up correctly, can be blank. Hostname verification for Kafka to Kafka connections. |
| <code>## The number of stream threads</code> | |

| Setting | Notes |
|---|---|
| <code>predeploy.stream.num.threads=6</code> | Do not change unless performance issue. |
| ## Log level for each EB container | |
| <code>predeploy.level=info</code> | Support settings only. |
| <code>predeploy.kafka.log.level=\${predeploy.level}</code> | |
| <code>predeploy.zookeeper.log.level=\${predeploy.level}</code> | |
| <code>predeploy.schema.log.level=\${predeploy.level}</code> | |
| <code>predeploy.web.service.log.level=\${predeploy.level}</code> | |
| <code>predeploy.c2av.stream.processor.log.level=\${predeploy.level}</code> | |
| <code>predeploy.eventbroker.routing.processor.log.level=\${predeploy.level}</code> | |
| ## Host path directory for ArcMC certificates | |
| <code>predeploy.arcmc.certs.path=/opt/arcsight/k8s-hostpath-volume/eb/arcmccerts</code> | |
| ##truncate fields in c2av | |
| <code>predeploy.c2av.field.truncate=false</code> | <p>If true, fields that are too long will be truncated to fit in the SuperSchema. See ArcMC Admin Guide for details of SuperSchema.</p> <p>If false (default), data in large fields will not be searchable.</p> |

Glossary

A

Apache Avro

A data serialization system. Avro enables highly space-efficient event storage.

Apache Flume

A service for efficiently collecting, aggregating, and moving large amounts of log data.

Apache Hadoop

A software framework that enables the distributed processing of large data sets across clusters of computers. It is designed to scale up from single servers to thousands of machines, each offering local computation and storage. You can configure Hadoop as an Event Broker consumer.

Apache Kafka

An open source distributed publish-subscribe messaging system installed as part of Event Broker.

Apache ZooKeeper

A centralized service for maintaining configuration information, naming, providing distributed synchronization, and providing group services. ZooKeeper's architecture supports high availability through redundant services. ZooKeeper is installed as part of Event Broker, which uses it to coordinate the Kafka cluster.

ArcMC

ArcSight Management Center (ArcMC) is an ArcSight product that....

B

broker

An instance of the Kafka server software.

C

CEF

Common Event Format (CEF) is an extensible, text-based, high-performance format designed to support multiple device types in the simplest manner possible. Various message syntaxes are reduced to one-matching ArcSight Enterprise Security Manager (ESM) normalization. Specifically, CEF defines a syntax for log records comprised of a standard header and a variable extension, formatted as key-value pairs. This format contains the most relevant event information, making it easy for event consumers to parse and use them.

channel

In Apache Flume, a buffer that stores events, until a sink has successfully written the them.

cluster

A collection of brokers working together to increase throughput and durability.

consumer

A process that subscribes to one or more topics and processes the feed of messages.

consumer group

A logical grouping of several consumers, where only one consumer in the group will process each message.

consumer offset

The read position for a consumer in a partition.

D

device group

In Logger, a category of named source IP addresses called devices. Device groups can be associated with storage rules that define the storage group where events from specific devices are stored. Refer to your Logger documentation for complete details.

E

ESM

Enterprise Service Management (ESM) is an ArcSight product that...

Event Broker Manager

Administration tool packaged with Event Broker. Equivalent to Yahoo Kafka Manager.

F

FIPS

Federal Information Processing Standards (FIPS) are standards developed by the U.S government for use in computer systems by non-military government agencies and government contractors. Specifically, FIPS PUB 140-2, is a U.S. government computer security standard used to approve cryptographic modules.

H

HDFS

Hadoop Distributed File System (HDFS) is a Java-based file system that provides scalable and reliable data storage, and was designed to span large clusters of commodity servers.

I

installer.properties

Properties file that controls many Event Broker settings.

Investigate

Investigate is an ArcSight product that...

L

leader

The broker containing the original replica of a partition, and manages that data.

Logger

An ArcSight product that receives event data and stores it for retrieval and analysis. You can configure Logger as an Event Broker consumer. Refer the Logger Administrator's Guide for complete details.

N

node

The machine a Kafka instance is running on.

O

offset

A sequential number identifying the location of a message in a partition. The sum of partition offsets is the total number of events in the topic.

P

partition

A segment of a topic. There can be one or more partitions for each topic. The number of partitions limits the maximum number of consumers in a consumer group.

pool

A logical grouping of Loggers. Loggers in a pool belong to the same consumer group and subscribe to the same topics.

producer

A process that publishes messages to a topic. In Event Broker, this is an ArcSight SmartConnector.

publish

The action of sending topics to the Event Broker.
A producer publishes event on a given topic.

Q

quorum

The set of all in-sync replicas for a particular partition. Replicas are considered in-sync if they are caught-up to the leader. The leader waits until a majority of replicas have received the data before considering it to be committed. On leader failure, a new leader is elected through the coordination of a majority of the followers. If there is an odd number of replicas, a majority is ensured. Any replica in the quorum can become the leader. This enables the producer to continue to publish messages and the consumer continues to receive the correct messages, even when there is failure.

R

receiver

In Logger, the process that receives events, captures event data, and populate each event with information about its origin. Refer to the Logger Administrator's Guide for complete details.

replica

A copy of a partition. There can be one or more per partition; even if there is no redundancy, the original is still called a replica.

replication factor

The number of times a topic is duplicated across Kafka nodes. A replication factor of 3 means that the topic is copied to 3 Kafka nodes.

route

A rule that directs Event Broker to copy events that meets certain criteria to the route's destination topic. A route leaves a copy of the event in both the source and destination topics.

S

scheduler

Manages and tracks the job process for Kafka.

sink

In Apache Flume, the sink forwards events to the storage destination.

SmartConnector

An ArcSight product that collects event data from objects on your network. They normalize the data in two ways: normalizing values (such as severity, priority, and time zone) into a common format, and normalizing the data structure into a common schema. You can configure SmartConnectors as Event Broker producers. Refer to the SmartConnector User's Guide for complete details.

SOC

Security Operations Center

source

In Apache Flume, a source sends events to Flume.

Stream Processor

The Event Broker mechanism for processing incoming events and converting their data format from CEF to AVRO. Also known as c2av.

subscribe

The action a consumer takes in order to receive the events that are published to a topic. A subscriber can receive the events published while the subscriber is active, or it can request events "from the beginning of time" the first time its consumer group is seen. From then on it will retrieve events since the last time a consumer for its group retrieved events.

T

TLS

Transport Layer Security. Enabled in Event Broker by default.

topic

A feed of messages relating to the same category.

transform

To convert data from one format to another. For example, Event Broker transforms CEF events into Avro format, where they can be stored in Vertica.

V

Vertica

HPE Vertica is an advanced SQL database analytics portfolio that enables you to run SQL on Hadoop, and leverage scalable predictive analytics and a comprehensive library of built-in analytical functions.

Y

Yahoo Kafka Manager

An open source tool for managing Apache Kafka. Event Broker includes a version of Yahoo Kafka Manager

Send Documentation Feedback

If you have comments about this document, you can [contact the documentation team](#) by email. If an email client is configured on this system, click the link above and an email window opens with the following information in the subject line:

Feedback on Administrator's Guide (Event Broker 2.02)

Just add your feedback to the email and click send.

If no email client is available, copy the information above to a new message in a web mail client, and send your feedback to arc-doc@hpe.com.

We appreciate your feedback!