
Micro Focus Security ArcSight Event Broker

Software Version: 2.20

Deployment Guide

Document Release Date: April 15, 2018

Software Release Date: April 2018



Legal Notices

Warranty

The only warranties for products and services of Micro Focus and its affiliates and licensors ("Micro Focus") are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Micro Focus shall not be liable for technical or editorial errors or omissions contained herein. The information contained herein is subject to change without notice.

Restricted Rights Legend

Confidential computer software. Except as specifically indicated otherwise, a valid license from Micro Focus is 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 2016-2018 Micro Focus or one of its affiliates.

Trademark Notices

Adobe™ is a trademark of Adobe Systems Incorporated.

Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

Support

Contact Information

Phone	A list of phone numbers is available on the Technical Support Page: https://softwaresupport.softwaregrp.com/support-contact-information
Support Web Site	https://softwaresupport.softwaregrp.com/
ArcSight Product Documentation	https://community.softwaregrp.com/t5/ArcSight-Product-Documentation/ct-p/productdocs

Revision History

Date	Description
April 15, 2018	Initial release of this document.

Contents

Chapter 1: Overview	5
About ArcSight Event Broker	5
Deployment Process Outline	5
Plan the Deployment Topology	6
Single Master Deployment Topology	7
Multi-Master Deployment Topology	8
Chapter 2: Prepare Systems for Deployment	10
System Requirements	10
Set up the Event Broker Nodes (Master and Worker)	11
Prepare Logical Volumes on Cluster Nodes (Optional)	11
Plan the Cluster's Security Configuration	12
Prepare a Virtual IP for High Availability (multi-master only)	13
Configure the Network Settings	14
Configure the Firewall Settings	14
Configure an NFS Server	15
Internal NFS	15
External NFS Server	16
Configure Proxy Settings	18
Increase the per-user process limits	18
Configure Network Time Protocol	18
Activate your Docker Hub Account (Online Deployment Only)	19
Chapter 3: Install the ArcSight Installer and Event Broker	20
Install the ArcSight Installer	20
Perform Pre-Deployment Configuration	21
Load images to the Local Docker Registry (Online Method)	22
Load images to the Local Docker Registry (Offline Method)	23
Deploy Product Images	24
Perform Post-Deployment Configuration	24
Install your license before the evaluation period ends	25
Chapter 4: Uninstall Event Broker and ArcSight Installer	26
Chapter 5: Upgrade to Event Broker 2.20	27
Appendix A: Troubleshooting	29
Appendix B: Installer Command Line Arguments	30
Appendix C: The arcsight-installer.properties file	32

Send Documentation Feedback	37
-----------------------------------	----

Chapter 1: Overview

This document describes how to deploy ArcSight Event Broker using the ArcSight Installer.

About ArcSight Event Broker

ArcSight Event Broker centralizes event processing and enables event routing, which helps you to scale your ArcSight environment, and opens ArcSight event data to third-party solutions. Event Broker enables you to take advantage of scalable, highly available, multi-broker clusters for publishing and subscribing to event data. Event Broker integrates with ADP ArcSight SmartConnectors and Collectors, Logger, ESM, and Investigate; can be managed and monitored by ArcSight Management Center (ArcMC); and is foundational for using ArcSight ADP products.

- After you install and configure Event Broker you can use ADP ArcSight SmartConnector and Collectors and to publish data, and subscribe to that data with ADP Logger, ArcSight ESM, Apache Hadoop, or your own custom consumer.
- Event Broker supports both Common Event Format (CEF) 0.1 and 1.0.
 - CEF 0.1 is the legacy ArcSight CEF version that supports IPv4 addresses available with SmartConnector version 7.4 and earlier.
 - CEF 1.0, available with SmartConnector version 7.5 and later and Collectors version 7.8 and later, supports IPv4 and IPv6 addresses. Note, however, that although it supports IPv6 event content, Event Broker does not support installation on IPv6 only systems.
- Event Broker provides a packaged version of Apache Kafka 1.0.0.
- Event Broker manages the distribution of events in topics to which consumers can subscribe. There are three default topics you can configure as your destinations:
 - eb-cef: accepts CEF event data.
 - eb-esm: accepts binary security events, which is the format consumed by ArcSight ESM.
 - eb-con-syslog: if you are using the Connector in Event Broker (CEB) feature, you can send syslog data using a Collector to this topic.

In addition, you can create new custom topics to which your SmartConnectors and Collectors can connect.

Event Broker features and functionality are explained in detail in the Event Broker Administrator's Guide, available from the ArcSight documentation community on [Protect724](#).

Deployment Process Outline

The complete process of deploying Event Broker includes these steps:

1. **Deployment Preparation.** Prepare and provision your network and dedicated hosts.
2. **Install the ArcSight Installer.** The ArcSight Installer is the platform used to install ArcSight Event Broker and Investigate.
3. **Deploy Event Broker.** You can deploy Event Broker using [online deployment](#), where you download the product images from the online Docker Hub, or [offline deployment](#), where you download the product images from the ArcSight software entitlement site.
4. **Licensing:** Install your permanent license, to ensure continuity of functionality and event flow.

Each of these steps is explained in the following sections.

For detailed instructions on the operation and management of Event Broker after initial deployment, see the Event Broker Administration Guide, available from the ArcSight support community at [Protect724](#).

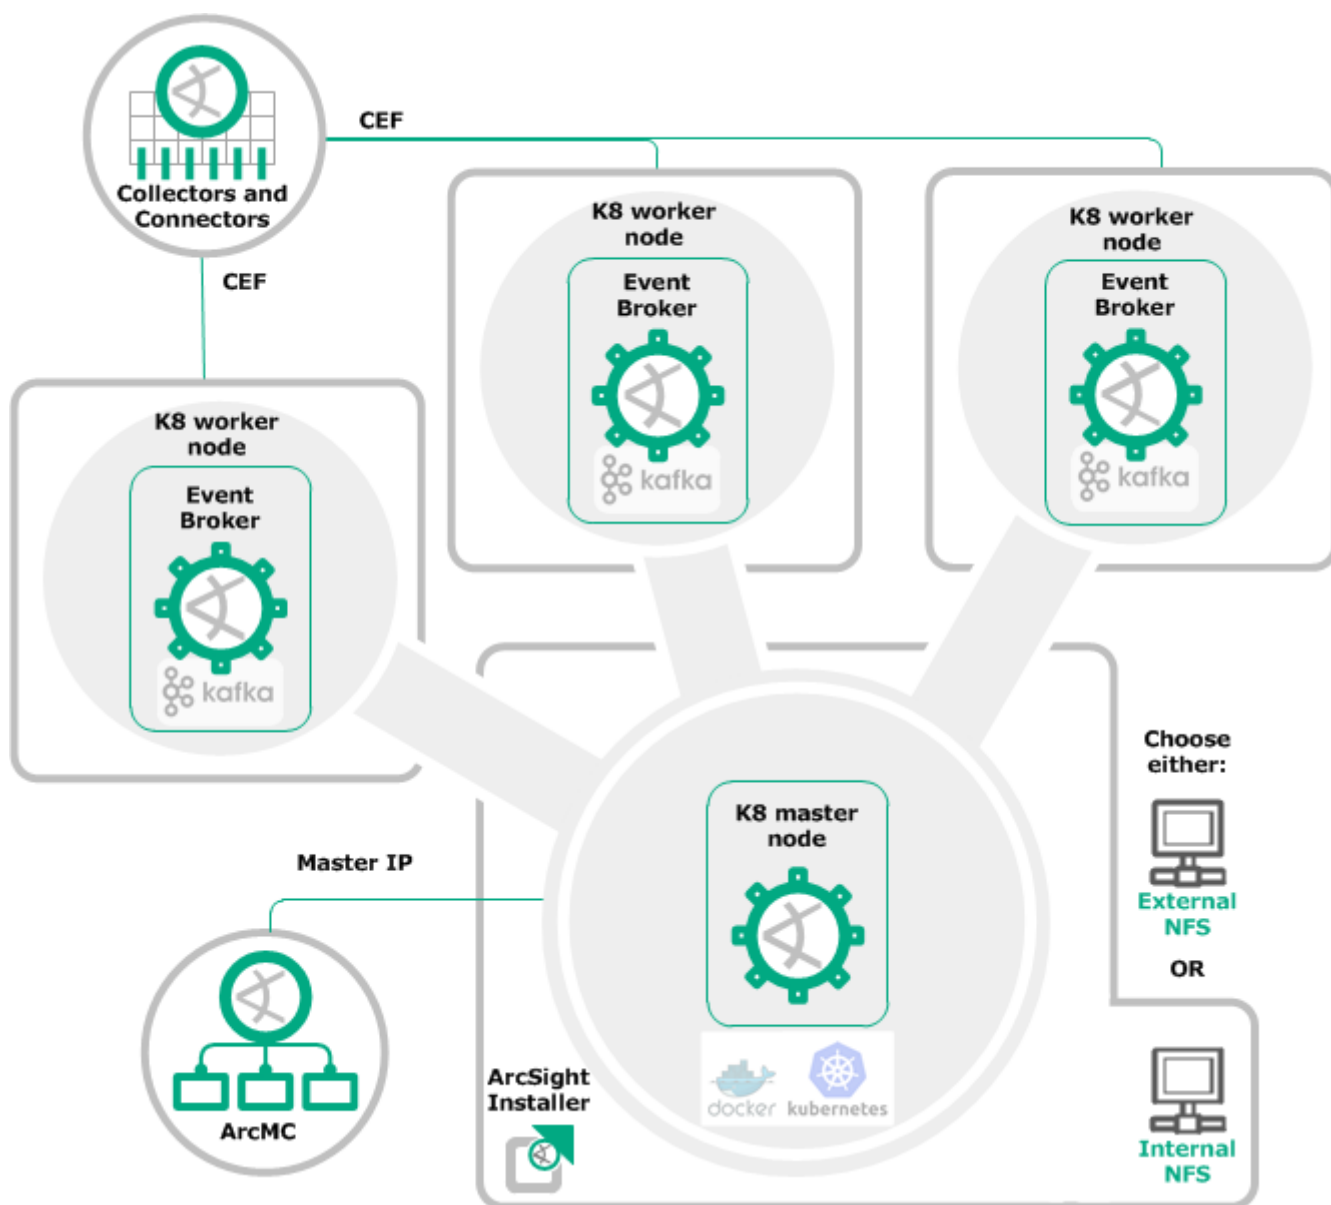
Plan the Deployment Topology

You can choose from multiple options when planning your deployment topology. Each option has benefits and disadvantages.

Single master: This topology requires fewer master nodes. It does not provide high availability or failover of the master node.

Multi-master: This topology requires more master nodes. It provides high availability and failover. This enables the cluster to continue function if one master node fails.

Single Master Deployment Topology

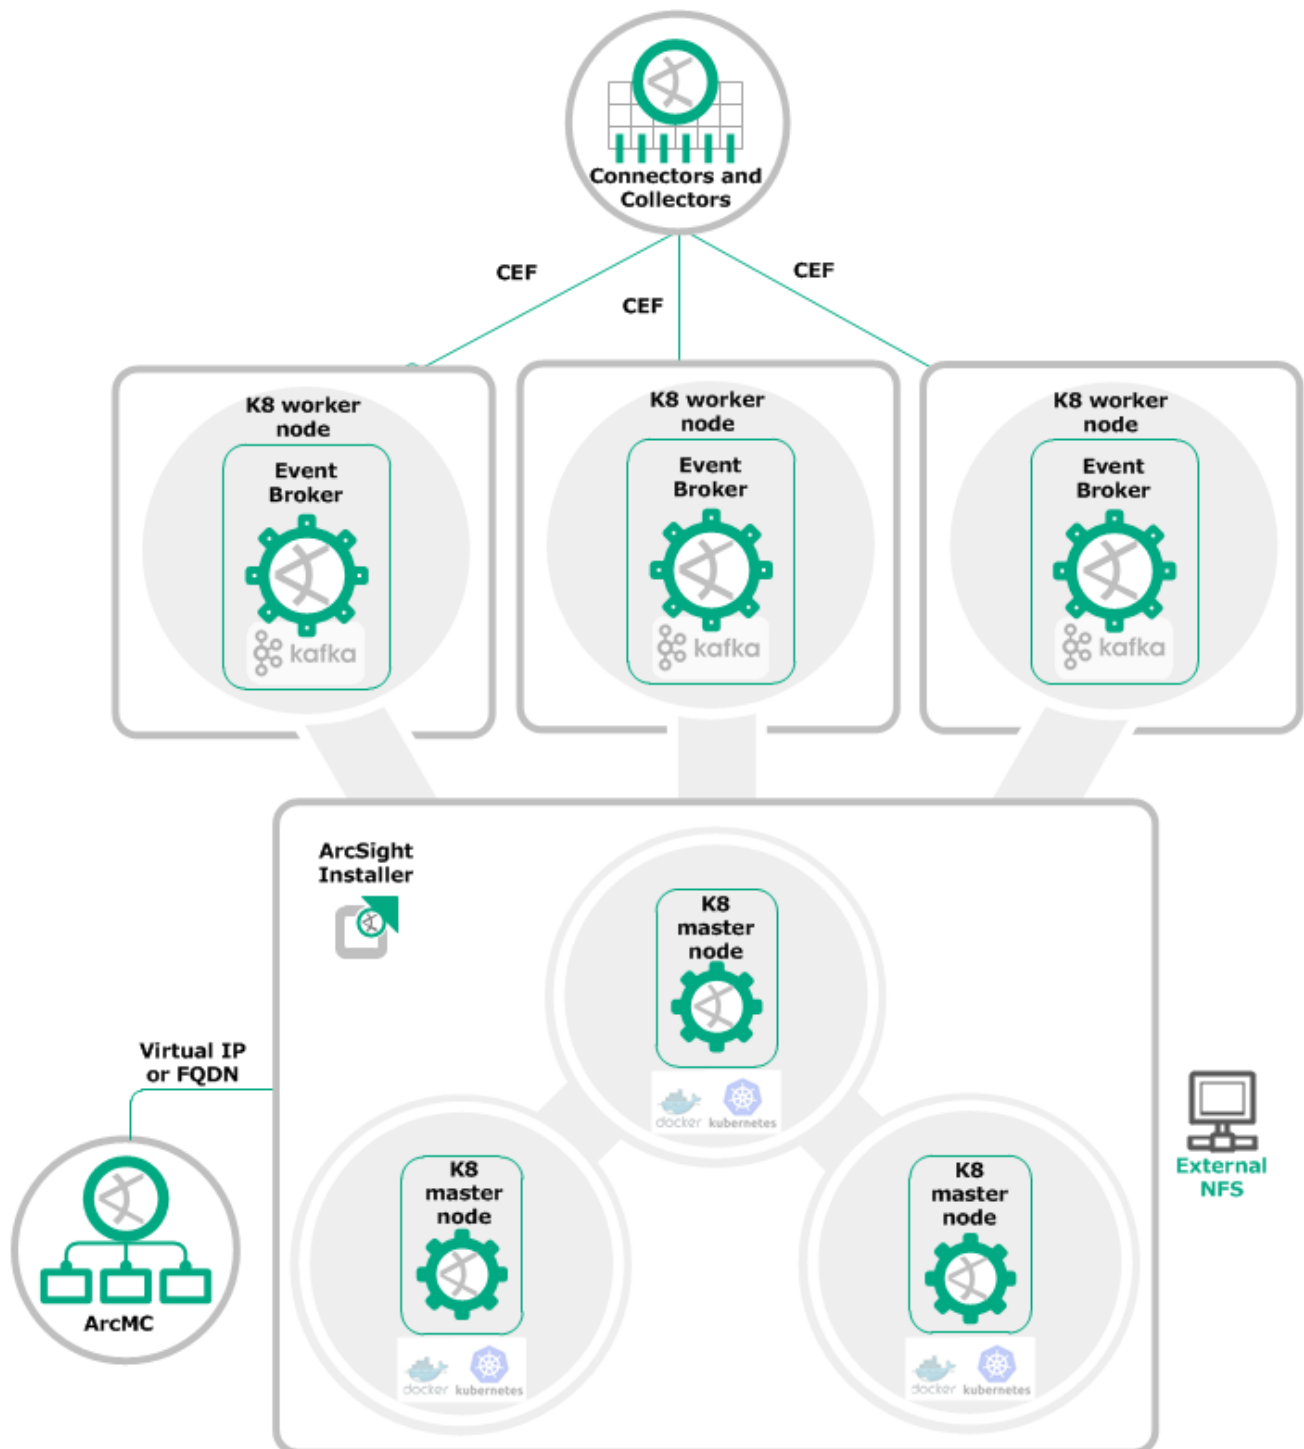


In a single master cluster deployment, there are 4 or more nodes in the cluster.

- Master node: One dedicated master node that runs ArcSight Installer and Event Broker Web Service container.
- Worker nodes: Three or more systems that run all other Event Broker pods including Kafka and Zookeeper.
- SmartConnector and Collectors connect directly to the worker nodes to send events.
- ArcMC connects to the master node.

- The cluster uses either an internal or external NFS server to store persisted data for ArcSight Installer, Event Broker, and Investigate if it is installed as well.
- This deployment topology does not provide high availability or failover of the master node.

Multi-Master Deployment Topology



In a multi-master cluster deployment, there are three master nodes and three or more worker nodes in the cluster.

- Master nodes: Three master nodes that host ArcSight Installer and Event Broker Web Service container. These nodes collectively provide high-availability and failover if one master node goes down. Application clients connect to the cluster using a Virtual IP address or FQDN. The request is directed to one of the master nodes. The FQDN is used when specifying the cluster host name in any configurations.
- Your network is configured with a Virtual IP address and FQDN that forwards requests to the systems identified as master nodes in the cluster.
- Worker nodes: Three or more systems that run Kafka and Zookeeper pods.
- Connectors and Collectors connect directly to the worker nodes to send events.
- ArcMC connects to the master nodes using the Virtual IP address or FQDN that one of the master nodes listens on.
- The cluster uses an external NFS share to store persisted data used by ArcSight Installer, Event Broker, and Investigate (if it is installed).

This deployment topology DOES provide high availability or failover of the master nodes. One master node can go down, and the cluster will continue to function. Two of the three master nodes must continue to run for a fully functional cluster. For production environments, it is strongly recommended that you choose the multi-master deployment topology.

Chapter 2: Prepare Systems for Deployment

This section discusses the following preparatory steps to take before deployment.

• System Requirements	10
• Set up the Event Broker Nodes (Master and Worker)	11
• Prepare Logical Volumes on Cluster Nodes (Optional)	11
• Plan the Cluster's Security Configuration	12
• Prepare a Virtual IP for High Availability (multi-master only)	13
• Configure the Network Settings	14
• Configure the Firewall Settings	14
• Configure an NFS Server	15
• Configure Proxy Settings	18
• Increase the per-user process limits	18
• Configure Network Time Protocol	18
• Activate your Docker Hub Account (Online Deployment Only)	19

System Requirements

Installation of the ArcSight Installer platform has the following system requirements.

- **Disk Space:** Required disk space for each master node and worker node is as follows:

Partition	Master Node	Worker Node	Notes
/opt	200 GB	150 GB	ArcSight Installer must be installed under the /opt directory.
/var	10 GB	10 GB	
/root	N/A	10 GB	

- **Operating System:** Supported operating systems are discussed in the ADP Support Matrix, available from [Protect724](#).
- **Node Sizing:** Information on Event Broker node sizing requirements can be found in the Event Broker Administrator's Guide.
- SELinux must be disabled or running in permissive mode.
- The file system type must be ext4.
- All commands in the installer process must run as a user with root privileges.
- The following packages need to be installed using `rpm/yum` on the master and all workers:

```
yum install -y unzip nfs-utils libseccomp libtool-ltdl
```

- Java 8 JRE (version 1.8, update 131 or later) must be present on the master node. Update or install the package using yum or rpm to ensure the minimum version is installed.

Set up the Event Broker Nodes (Master and Worker)

Provision your Master and Worker Nodes

Note: Previous versions of Event Broker supported a single-master installation (that is, 1 master node). This installation model is still supported, but will not provide the same level of availability and failover as will a multi-master (3 master node) cluster.

Set up the systems in the cluster

Master Nodes: Depending on your chosen deployment topology, single master or multi-master, set up the systems that will function as the master node. Make sure that each system meets the minimum [system requirements](#) for a master node. The recommended configuration for a production Event Broker is 3 master nodes.

Worker Nodes: set up the systems that will function as the worker nodes. Make sure that each system meets the minimum [system requirements](#) for a worker node.

Enable Password-less SSH on all systems in the cluster

To enable the master nodes to communicate with worker nodes over SSH during installation (or upgrade), generate a key pair on any master node and copy the public key to every other master and worker node.

To generate the key pair, run the following command on the master node where you will execute the ArcSight Installer installation script:

```
ssh-keygen -t rsa
```

On the same master node, run the following command to copy the generated public key to all other master and worker nodes in the cluster:

```
ssh-copy-id -i ~/.ssh/id_rsa.pub root@<node_fqdn_or_ip>
```

Note: SSH access only needs to be enabled for installation (or upgrade). If required by your security policy, you can disable SSH access after installation or upgrade operations are complete.

Prepare Logical Volumes on Cluster Nodes (Optional)

Follow the steps below on each cluster node to ensure that you have enough logical volumes for the ArcSight Installer platform installation. You can choose any volume group name, logical volume name

and disk location address for your installation, according to your system.

Note: This is an optional step.

Prepare a physical disk for the Kubernetes cluster nodes. The physical host of your system must meet the [system requirements](#).

1. Create a volume group by running the following command:

```
vgcreate [volume group name] [logical volume name]
```

For example: `vgcreate core-platform /dev/sdb`

2. Create a logical volume for the ArcSight Installer installation by running the following command:

```
lvcreate -l 100%FREE -n [logical volume name] [volume group name]
```

For example, to utilize 100% of the volume group:

```
lvcreate -l 100%FREE -n mylv core-platform
```

3. Activate the volume group by running the following command:

```
vgchange -ay [volume group name]
```

For example:

```
vgchange -ay core-platform
```

4. Format the file system to ext4 by running the following command:

```
mkfs.ext4 [logical volume path]
```

For example: `mkfs.ext4 /dev/core-platform/mylv`

5. Mount the volumes under the folder in which you will install the ArcSight Installer by running the following command:

```
mount [logical volume path] [platform installation folder]
```

For example:

```
# mount /dev/core-platform/mylv /opt/arcsight/kubernetes
```

Plan the Cluster's Security Configuration

Before installing Event Broker, determine which security mode you want for communication between ArcSight components. The security mode of consumers and producers connected to Event Broker must be the same as that set for Event Broker. Set up the other ArcSight components with the security mode you intend to use before connecting them to the Event Broker. Changing the security mode after Event

Broker has been deployed will require system down time. If you do need to change the security mode after Event Broker deployment, see the *Event Broker Administrator's Guide*.

Product	Preparations needed	Open ports	Supported security modes	Where to find more information
ArcMC	Install ArcMC before Event Broker installation.	38080	<ul style="list-style-type: none"> • TLS • FIPS • ClientAuth 	<i>ArcMC Administrator's Guide</i>
ArcSight SmartConnectors and Collectors	<p>ArcSight SmartConnectors and ArcMC onboard connectors can be installed and running prior to installing Event Broker.</p> <p>FIPS mode setup is not supported between SmartConnector version 7.5 and Event Broker. TLS and ClientAuth are the only security modes supported between SmartConnector version 7.5 and Event Broker. FIPS mode is supported between Connectors version 7.6 and above and Event Broker.</p>	9093	<ul style="list-style-type: none"> • TLS • FIPS • ClientAuth 	<i>SmartConnector User Guide</i> <i>ArcMC Administrator's Guide</i>
ArcSight ESM	ArcSight ESM can be installed and running prior to installing Event Broker. Note that changing ESM from FIPS to TLS mode requires a redeployment of ESM; see the ESM documentation for more information.	9093	<ul style="list-style-type: none"> • TLS • FIPS • ClientAuth 	<i>ESM Installation Guide</i> <i>ESM Administrator's Guide</i>
ArcSight Logger	ArcSight Logger can be installed and running prior to installing Event Broker.	9093	<ul style="list-style-type: none"> • TLS • FIPS • ClientAuth 	<i>Logger Administrator's Guide</i>

Prepare a Virtual IP for High Availability (multi-master only)

A virtual IP address (VIP) is an IP address shared by all master nodes in a cluster deployed in a multi-master configuration. The VIP is used for redundancy by providing failover in a multi-master configuration. When a master node goes down, another master node takes over the VIP and responds to requests sent to the VIP.

Ask your network administrator to configure a VIP as follows:

- allocate a free IP address in the same subnet as the master(s)
- create a fully qualified domain name (FQDN), and bind the IP address to the FQDN.

You will use the IP address as the HA VIP and the FQDN as your Kubernetes cluster host name.

In a multi-master cluster deployment, you will launch the applications (for example, ArcSight Installer or Investigate) using the Virtual IP, rather than the master node IP. The port remains the same (port 5443 for ArcSight Installer and default port 80 for Investigate).

Configure the Network Settings

To prepare your network for the deployment process, take the following steps.

- Ensure host name resolution through DNS across all nodes in the cluster, including correct forward and reverse DNS lookups. Host name resolution must not be performed through `/etc/hosts`.
- Ensure that all master and worker nodes are configured with an FQDN (fully-qualified domain name), and are in the same subnet.

Note: Event Broker does not support installation on IPv6-only networks. However, it does support ingestion of event data that contains both IPv4 and IPv6 address.

- Event Broker uses the host system FQDN for Kafka advertised `.host.name`. If FQDN resolves in the NAT environment, then producers and consumers will function correctly. If there are network-specific issues resolving FQDN through NAT, then DNS will need to be updated.
- If you are performing an [online deployment](#), enable internet access in order to download the product container images.

Configure the Firewall Settings

Make sure that the `firewalld.service` is enabled and running on all nodes before running the `arcsight-installer-master.sh` and `arcsight-installer-add-node.sh` scripts. The following firewall ports will be opened during the installation process.

Used by	Port
ArcSight Installer	5443
Kubernetes	2379, 2380, 3000, 4001, 4194, 5000, 8080, 8088, 8200, 8285, 8443, 10248-10252, 10255
NFS	111, 2049, 20048, 37189
Event Broker	2181, 9092, 9093, 38080, 39000, 39093, 32181
CEB	39001-39010
CAdvisor	4194

The NFS ports are used only in clusters that are configured to use an internal NFS server.

Avoiding Possible Conflicts with Network Ranges

The ArcSight Installer configures firewall settings during setup (assuming `firewalld.service` is up and running) on both Kubernetes master and worker nodes.

The Installer will use the following network ranges by defaults:

- 172.16.0.0/16: Subnetwork of 65,536 addresses for operation of Kubernetes pods with containers running in them. Each pod will operate with /24 subnetwork from the following range.
- 172.30.78.0/24: Subnetwork of 256 addresses for operation of Kubernetes services, including internal Kubernetes DNS service, located on pod 172.30.78.78.

For best results, make sure your network is conflict-free with the /16 and /24 ranges of addresses. If those are occupied or inaccessible due to network configuration, make sure to utilize another range by making corresponding changes to `POD_CIDR`, `SERVICE_CIDR` and `DNS_SVC_IP` parameters in the `arcsight-installer-master.sh` script before executing it.

Configure an NFS Server

The ArcSight Installer platform and ArcSight products require an NFS server to operate. You can choose to use the default internal NFS server, or to use an external NFS server.

Internal NFS

The ArcSight Installer can create a default internal NFS server on the master node for shared use by the ArcSight Installer and products. To configure this during the installation, you will use the argument `--NFS_SERVER=internal` when you install the ArcSight Installer.

Note: Use of the default internal NFS server is only recommended for single-node and non-production deployment environments. For a production environment with multiple master nodes, use an external NFS server.

Some security hardening is performed by the installer, but it is strongly recommended that you take additional hardening actions, such as adding firewall rules restricting access only from the master and worker subnet for the following services:

- NFS Server (on port 2049/tcp and 2049/udp)
- rpcbind (on port 111/tcp)
- rpc.mountd (on port 20048/tcp)

External NFS Server

In the case of a multi-master deployment (3 master nodes), NFS should run on an external server which is highly available.

The following setup procedure assumes:

- 3 master nodes with IP addresses: 10.1.2.11 - 13
- 3 worker nodes with IP addresses: 10.1.2.21 - 23
- Event Broker cluster subnet for Kubernetes pods (POD_CIDR): 172.16.0.0/16
- An NFS share root directory of /opt/arcsight/nfs/volumes

To set up the External NFS share:

1. Log in to the external NFS server as root or as a sudo user.
2. Run the command `rpm -qa | grep rpcbind` to make sure that the `rpcbind` package is installed on the host. If the package is not already installed, run the following command to install it: `yum install rpcbind`
3. Run the following command to install the NFS server: `yum install -y nfs-utils`
4. Run the following commands to enable the `rpcbind` and `nfs-server` services:


```
systemctl enable rpcbind
systemctl start rpcbind
systemctl enable nfs-server
systemctl start nfs-server
```
5. On the NFS server, run the following commands to create NFS share directories for Event Broker and Investigate. If you do not plan to deploy Investigate on this cluster, do not run the commands specific to Investigate.


```
mkdir -p /opt/arcsight/nfs/volumes/itom/core
mkdir -p /opt/arcsight/nfs/volumes/eventbroker
mkdir -p /opt/arcsight/nfs/volumes/investigate
```
6. Add a group (GID=1999) and user (UID=1999):


```
groupadd -g 1999 eventbroker
useradd -g 1999 -u 1999 eventbroker
```

Note: Ignore this step if the GID and UID already exist.
7. Set the owner/group for the NFS share directories:


```
chown -R 1999:1999 /opt/arcsight
```
8. In the `/etc/exports` file, add the following lines:


```
/opt/arcsight/nfs/volumes/itom/core 10.1.2.11(rw,sync,anonuid=1999,anongid=1999,all_squash)
10.1.2.12(rw,sync,anonuid=1999,anongid=1999,all_squash) 10.1.2.13
```



```
(rw,sync,anonuid=1999,anongid=1999,all_squash) 10.1.2.21
(rw,sync,anonuid=1999,anongid=1999,all_squash) 10.1.2.22
(rw,sync,anonuid=1999,anongid=1999,all_squash) 10.1.2.23
(rw,sync,anonuid=1999,anongid=1999,all_squash) 172.16.0.0/16
(rw,sync,anonuid=1999,anongid=1999,all_squash)
/opt/arcsight/nfs/volumes/eventbroker 10.1.2.11(rw,sync,anonuid=1999,anongid=1999,all_squash)
10.1.2.12(rw,sync,anonuid=1999,anongid=1999,all_squash) 10.1.2.13
(rw,sync,anonuid=1999,anongid=1999,all_squash) 10.1.2.21
(rw,sync,anonuid=1999,anongid=1999,all_squash) 10.1.2.22
(rw,sync,anonuid=1999,anongid=1999,all_squash) 10.1.2.23
(rw,sync,anonuid=1999,anongid=1999,all_squash) 172.16.0.0/16
(rw,sync,anonuid=1999,anongid=1999,all_squash)
/opt/arcsight/nfs/volumes/investigate 10.1.2.11(rw,sync,anonuid=1999,anongid=1999,all_squash)
10.1.2.12(rw,sync,anonuid=1999,anongid=1999,all_squash) 10.1.2.13
(rw,sync,anonuid=1999,anongid=1999,all_squash) 10.1.2.21
(rw,sync,anonuid=1999,anongid=1999,all_squash) 10.1.2.22
(rw,sync,anonuid=1999,anongid=1999,all_squash) 10.1.2.23
(rw,sync,anonuid=1999,anongid=1999,all_squash) 172.16.0.0/16
(rw,sync,anonuid=1999,anongid=1999,all_squash)
```

9. Export the newly added directories:

```
exportfs -ra
```

NOTE: As an option you can replace steps #6, 7, and 8 with the following:

1. Copy the setupNFS.sh script from the initial master node to the NFS server. This file is located in the /opt/arcsight/kubernetes/scripts directory on the initial master node.
2. Run the setupNFS.sh script once for each directory share created

```
sh setupNFS.sh /opt/arcsight/nfs/volumes/itom/core
sh setupNFS.sh /opt/arcsight/nfs/volumes/eventbroker
sh setupNFS.sh /opt/arcsight/nfs/volumes/investigate
```

IMPORTANT: During ArcSight Installer installation, the following arguments must be passed to the arcsight-installer-master.sh script:

```
--NFS_SERVER=<nfs_server_IP_address_or_hostname>
--NFS_FOLDER_ROOT=<root_nfs_folder>
```

where:

- NFS_SERVER is the IP address or hostname of the external NFS server
- NFS_FOLDER_ROOT is the root directory for the NFS shares used by Event Broker, such as /opt/arcsight/nfs/volumes in the example above.

Configure Proxy Settings

If a connection with the Internet is needed (such as for an online installation, to connect to Docker Hub), and if your organization requires it, you may have to specify a proxy server for http and https connections. On each cluster node, set the proxy settings by editing the `~/ .bashrc` file.

```
# vi ~/.bashrc

export http_proxy=<address of proxy server>

export HTTP_PROXY=<address of proxy server>

export https_proxy=<address of proxy server>

export HTTPS_PROXY=<address of proxy server>
```

If you have the `http_proxy` or `https_proxy` set, then `no_proxy` and `NO_PROXY` must contain at least the following values:

```
no_proxy=localhost, 127.0.0.1, <all cluster node IP addresses>,<all cluster
node FQDNs>,
<HA virtual IP Address>,<FQDN for the HA Virtual IP address>
```

Increase the per-user process limits

Perform the following steps on every master and worker node.

1. Open (or create, if necessary) the file `/etc/security/limits.d/20-nproc.conf`.
2. Add the lines below to the file, including the leading asterisks.


```
* soft nproc 10240
* hard nproc 10240
* soft nofile 65536
* hard nofile 65536
* soft core unlimited
* hard core unlimited
```
3. Reboot all master and worker nodes. Nodes can be rebooted in any order.
4. Verify that all nodes are up and running by running the following command.


```
ulimit -a
```

Configure Network Time Protocol

`chrony` is a versatile implementation of the Network Time Protocol (NTP) and keeps the system clock of each cluster node in sync. A network time server must be available. `chrony` is installed by default on

some versions of RHEL/CentOS. Verify `chrony` configuration by using the command:

```
chronyc tracking
```

If `chrony` is not installed on any of the systems, install it with the following procedure.

1. `yum install chrony`
2. Update `/etc/chrony.conf` with the time server information

Start `chronyd` to start and enable the `chrony` daemon.

```
systemctl start chronyd
```

```
systemctl enable chronyd
```

3. Verify that `chrony` is operating correctly.

```
chronyc tracking
```

Activate your Docker Hub Account (Online Deployment Only)

Docker Hub is cloud-based registry service which can store manually pushed images, including ArcSight Event Broker product images.

To complete an online deployment, you must use or create a valid company Docker ID (hub.docker.com) to grant you instant and secure cloud access to Event Broker software. If you do not have an existing Docker account, please follow the steps below to register a Docker ID.

Note: If you plan to perform an offline deployment, you will download images from the ArcSight software entitlement site, and this step is not necessary.

1. Go to <https://hub.docker.com>
2. Create a Docker ID, enter your company email address, and create a password.
3. Click **Sign Up**.
4. Click the **Confirm Your Email** link in the email you received from Docker to confirm your Docker ID account.
5. Go to <https://hub.docker.com> to verify that you can log into Docker Hub.
6. After login, click your Docker ID on the top right of the page. Click **Settings** and take a screenshot to include your Docker ID and the linked email address.

Please email your corresponding regional contact below with your registered Docker ID as well as the screenshot for us to enable your Docker Account. Based on your region and existing entitlements, contact your licensing team to enable your Docker ID:

- For the Americas region, contact dockersupport.ams@hpe.com
- For the APJ region, contact dockersupport.apj@hpe.com
- For the EMEA region, contact dockersupport.emea@hpe.com

Chapter 3: Install the ArcSight Installer and Event Broker

Once your planning is complete and prerequisites are met, you are ready to install the ArcSight Installer and then deploy Event Broker. The following topics are discussed in this section.

• Install the ArcSight Installer	20
• Perform Pre-Deployment Configuration	21
• Load images to the Local Docker Registry (Online Method)	22
• Load images to the Local Docker Registry (Offline Method)	23
• Deploy Product Images	24
• Perform Post-Deployment Configuration	24
• Install your license before the evaluation period ends	25

Install the ArcSight Installer

The ArcSight Installer enables the deployment of ArcSight products, such as Event Broker.

1. Log in to the one of the master nodes as the root user. In this document, this system will be referred to as the *initial master node*.
2. If you are not logged in as the root user, sudo to root.
3. Download the installation file from the [ArcSight software entitlements site](#).
4. Unzip the installation file to the /opt directory on the initial master node.
5. Change into the /opt/arcsight-installer-1.40.13 directory.
6. Run the ./arcsight-installer-master.sh script to start the installation. Depending on the cluster topology, the command you execute will be different.

For a multi-master installation with an external NFS server, install the initial master node with this command:

```
./arcsight-installer-master.sh --MASTER_NODES="<initial_master_ipv4><one_space><master2_ipv4><one_space><master3_ipv4>" --HA_VIRTUAL_IP=<virtual_ip> --EXTERNAL_ACCESS_HOST=<fqdn_of_virtual_ip> --NFS_SERVER=<nfs_ip_addr> --NFS_FOLDER_ROOT=<root_nfs_folder>
```

If you are configuring a single master cluster with an internal NFS configuration, run the following command.

```
./arcsight-installer-master.sh --NFS_SERVER=internal
```

If you are configuring a single master cluster with an external NFS configuration, run the following command.

```
./arcsight-installer-master.sh --NFS_SERVER=<nfs_ip_addr> --NFS_FOLDER_ROOT=<root_nfs_folder>
```

Note: Values for each of these arguments are discussed earlier in the [Deployment Planning](#) chapter. For a complete list of installer arguments, see [Installer Command Line Arguments](#).

7. From the initial master node, install the additional master and worker nodes using the following example command. You will need to execute this script for each host in the cluster.

This command installs a master node if the <IPv4_node_address> value is listed in the --MASTER_NODES argument passed to the `arcsight-installer-master.sh` script. If the <IPv4_node_address> value is not in that list, the command installs a worker node.

```
cd /opt/arcsight/kubernetes/scripts
./arcsight-installer-add-node.sh <IPv4_node_address>
```

Note: If `arcsight-installer-add-node.sh` fails during installation of a new node, re-run the `arcsight-installer-add-node.sh` for that same node once more. If it fails again, the failed node must be removed, before the script can be re-run again.

Run the command `kubectl get nodes`. If the command shows the failed node in any other state than Ready, remove it using the command `kubectl delete node <Failed_Node_IP>`. You can then rerun the `arcsight-installer-add-node.sh` script.

Log in to the ArcSight Installer Web Application

1. Browse ArcSight Installer
 - multi-master cluster: `https://<virtual IP_or_fqdn>:5443`
 - single-master cluster: `https://<master_node_IP_or_fqdn>:5443`
2. Enter the default credentials *admin/cloud*. After the first successful log on, you will be forced to change the admin password to a non-default value.
3. In ArcSight Installer, on the **Node Management** page, make sure that all nodes are listed and have the status *READY*.

Perform Pre-Deployment Configuration

The following steps must be performed before you deploy Event Broker containers. Connect to the initial master node to perform the steps in each section.

Edit the `arcsight-installer.properties`

You will likely need to change the default configuration for certain product capabilities before deploying Event Broker containers (such as FIPS, Client Authentication, the data retention size and period for

messages, etc).

1. Open the `/opt/arcsight/installer/arcsight-installer.properties` file in a text editor, and edit the values accordingly.
2. After changing any values in the file, you must run the script `update-arcsight-installer-properties.sh` for the changes to take effect.

See ["The arcsight-installer.properties file " on page 32](#) for more information.

Label Nodes in the Cluster

You will need to label each worker node to identify which node the zookeeper and kafka pods will run on when deployed in the cluster. The following steps must be performed from one of the master nodes. You will apply the labels to worker nodes.

The number of the Kafka and ZooKeeper nodes that are labeled must be the greater than the values defined for `eb-kafka-count` and `eb-zookeeper-count` properties in the [/opt/arcsight/installer/arcsight-installer.properties](#) file. The default value for both properties is 3.

- The number of Zookeeper nodes must be odd to support high availability and failover.
 - The number of Kafka nodes can be even or odd, and must be at least 3. The number of Kafka nodes in the cluster will depend on the cluster design and how it optimizes data throughput, fault-tolerance, and other factors.
1. To label a worker node for Kafka, run `kubectl label --overwrite node {worker_node_ip} kafka=yes`
 2. To label a worker node for Zookeeper, run `kubectl label --overwrite node {worker_node_ip} zk=yes`
If the command returns a refused or timed-out connection, temporarily remove your proxy settings using `unset http_proxy` before repeating the `kubectl` command.
 3. Check that nodes are labeled correctly by running these commands:
`kubectl get nodes -L=kafka`
`kubectl get nodes -L=zk`
 4. Launch the ArcSight Installer web application, then navigate to the **Node Management** page. Make sure that all nodes are listed, have the correct labels, and have the status **READY**

You may now proceed to load product images to the local Docker Registry. Choose one of the following options to complete this next step:

- Online method: you download product images using your activated Docker Hub account.
- Offline method: you download product images from the ArcSight software entitlement site.

Load images to the Local Docker Registry (Online Method)

Follow these steps downloading the product images from Docker.com, and then upload them to the local docker registry. Ensure you have [activated your Docker account](#) before proceeding.

1. Connect to the initial master node.
2. Run the following command to download images locally:


```
cd /opt/arcsight/kubernetes/scripts
./downloadimages.sh --suite eventbroker --registry docker --org
arcsightsecurity
```
3. Enter your Docker credentials. (See "[Activate your Docker Hub Account \(Online Deployment Only\)](#)" [on page 19](#) for details on how to create your Docker Hub account, if you have not already done so.)
4. Enter the product version that you want to download.
5. Once the images have been downloaded, run the following command to upload them to the local Docker registry. This step can take up to approximately 5 minutes to complete.


```
./uploadimages.sh --suite eventbroker
```

Load images to the Local Docker Registry (Offline Method)

Follow these steps to download the product images from the [ArcSight software entitlement site](#), and then upload them to the local docker registry.

Verifying the Download: Micro Focus provides a digital public key to enable you to verify that signed software you download from the software entitlement site is indeed from Micro Focus and has not been manipulated in any way by a third party. Visit the following site for information and instructions:
<https://h20392.www2.hpe.com/portal/swdepot/displayProductInfo.do?productNumber=HPLinuxCodeSigning>

1. Log on to the initial master node.
2. Download the Event Broker tar archive `arcsight-eventbroker-2.20.33.tar` from the Micro Focus [software entitlements site](#) to a secure location; for example, `/opt/arcsight/download/arcsight-eventbroker-2.20.33.tar`.
3. Verify the digital signature before unpacking the tar file.

Verify that the tar file MD5 checksum output matches the content of the md5 file. For example:

```
cd /opt/arcsight/download/
md5sum arcsight-eventbroker-2.20.33.tar
cat arcsight-eventbroker-2.20.33.md5
```

Both outputs should match.

4. Once verified, unpack the tar file.

```
cd /opt/arcsight/download
tar -xvf arcsight-eventbroker-2.20.33.tar
```

The `/opt/arcsight/download/eventbroker` folder will contain the product images. Verify that the folder contains a set of files.

5. Upload the images to the local registry as follows:

```
cd /opt/arcsight/kubernetes/scripts
./uploadimages.sh --suite eventbroker --dir
/opt/arcsight/download/eventbroker
```

Deploy Product Images

Once you have uploaded images to the local docker registry, you are ready to deploy ArcSight Event Broker.

1. In the ArcSight Installer web application, browse to the **Deployment** page. The list of products should be displayed with status *OFF*
2. In the Event Broker row, click **Deploy**.
3. In the version dialog, select 2.20 and then click **Deploy**. Deployment status will be changed to *IN PROGRESS*.
4. Once the product deployment is finished, the status will be changed to *DEPLOYED*. Please give the process some time to complete. (This can take a few minutes.)
A popup window will show the list of pods, with status and memory usage.
To check the pod status, click **Details**. Once all pods have status *RUNNING*, the product can be considered ready to configure and use.
5. The **Undeploy** button will remove a product and all its containers from Kubernetes. This process will take a similar amount of time as a deploying images, from 2-5 minutes.

IMPORTANT: Post-deployment configuration changes are lost when Event Broker is undeployed. Each time you redeploy Event Broker, make sure to re-configure any post-deployment properties for the cluster. This will be addressed in a future release.

Perform Post-Deployment Configuration

Configuration

To configure ArcSight Event Broker, go to **Configuration > ArcSight Event Broker**. Change the required configuration parameters and click **Save**.

After changing a product setting, one or more containers of the product will be restarted in the cluster. Depending on which pods need to restart, there may be a brief interruption to some running applications.

Install your license before the evaluation period ends

Event Broker ships with a 90-day instant-on evaluation license, which will enable functionality for 90 days after installation. In order for Event Broker to continue working past the initial evaluation period, you will need to apply the ADP ArcMC license to Event Broker.

IMPORTANT: *To ensure continuity of functionality and event flow, make sure you apply the proper license before the evaluation license has expired.*

For details on how to apply a new license file to Event Broker, see the Licensing chapter of the *Event Broker Administrator's Guide*.

Chapter 4: Uninstall Event Broker and ArcSight Installer

To remove a single node from the cluster

1. Open an ssh connection to the node you want to remove.
2. Run the following script: `/opt/arcsight/kubernetes/uninstall.sh`
3. Enter Y and N at the prompts.
4. Reboot each node after the script has finished executing.

To uninstall Event Broker and ArcSight Installer completely

When uninstalling Event Broker and ArcSight Installer from a cluster, the uninstall script should be run on worker nodes, first, before it is run on master nodes. This procedure will remove application components, but will not delete data stored on cluster nodes.

1. Stop all Collectors and Connectors from sending events to Event Broker and all consumers from reading events as well.
2. Connect to each node in the cluster, and then run the following command:
`/opt/arcsight/kubernetes/uninstall.sh`.
3. Enter Y and N for the prompts.
4. Reboot the node after the script has finished executing.

To clean up existing data stored on cluster nodes (Optional)

After uninstalling Event Broker and ArcSight Installer, product data (including the event data stored in Kafka topics) is still kept in multiple folders on each system. You can delete the product data if it is not planned for retention or future recovery.

1. Connect to each system in the cluster, all masters and workers, and then run the following command:

`rm -rf /opt/arcsight /opt/kubernetes /root/.kube`
2. If you set up an external NFS server, manually delete the shared folders on the external NFS server.

Chapter 5: Upgrade to Event Broker 2.20

This procedure will upgrade the ArcSight Installer from version 1.30 to 1.40, and then using the Upgrade step in the upgraded ArcSight Installer.

Note: Upgrading from a single master to a multi-master installation, or changing an internal NFS to an external NFS, are **not** supported by this process.

Upgrade the ArcSight Installer

As an operational note, Event Broker will still be able to process events while the Installer is upgraded. However, the ArcSight Installer *will* be down while the ArcSight Installer is upgraded.

IMPORTANT: If password-less SSH was disabled after the last installation, it must be re-enabled before starting the upgrade process.

1. From the Micro Focus [software entitlements site](#), download ArcSight Installer 1.40.
2. Copy the archive to the master node of your cluster (such as `/opt/arcsight/downloads/arcsight-installer-1.40.13.zip`).
3. Verify that the zip file MD5 checksum output matches the content of the md5 file. For example:

```
cd /opt/arcsight/downloads/
md5sum arcsight-installer-1.40.13.zip
cat arcsight-installer-1.40.13.md5
```

Both outputs should match.

4. Verify the digital signature.
5. Extract the zip file:


```
unzip arcsight-installer-1.40.13.zip
```
6. Change working directory to the extract folder, for example:


```
cd /arcsight-installer-1.40.13
```
7. Run the upgrade script with parameter `--patch`:


```
./upgrade.sh --patch
```
8. Wait for the changes to be fully propagated and check for completion by running:

```
kubectl get pods --all-namespaces
```

Note: Pay special attention to pods `arcsight-installer`, `suite-installer` and `suite-db`.

9. Browse to ArcSight Installer web application and verify that version 1.40.13 is displayed in bottom left corner.

Change Pre-deployment properties (Optional)

If you need to change any of the pre-configuration properties (such as FIPS, Client Authentication, and so on), do this before proceeding to the next section.

1. Open `/opt/arcsight/installer/arcsight-installer.properties` on the initial master, and update the values accordingly.
2. Then, propagate the updated values to the installer, run the following command:
`/opt/arcsight/installer/update-arcsight-installer-properties.sh`.

Upgrade Event Broker

During the upgrade process, the Event Broker will be down and non-operational. Plan accordingly before performing the upgrade.

1. Depending on whether you are performing an online or offline installation, select one of the following sections for the steps to download images and load them to the local docker registry.
 - ["Load images to the Local Docker Registry \(Offline Method\)" on page 23](#)
 - ["Load images to the Local Docker Registry \(Online Method\)" on page 22](#)
2. Launch the ArcSight Installer web application, and click **Upgrade** for Event Broker. Event Broker is upgraded to version 2.20.
3. If Event Broker is being managed by ArcMC, you must migrate the ArcMC certificate.
 - Locate the ArcMC certificate file. See the `arcmc-certs-path` property in the `arcsight-installer.properties` file for this location.
 - Open ArcSight Installer > Event Broker Configuration link > ArcMC Monitoring tab.
 - Copy and paste the contents of the ArcMC certificate file into the ArcMC Certificate field in ArcSight Installer.
 - Click the Save button.

It will may take up to 5 minutes for the configuration changes to propagate. To verify the status, check ArcMC under the Node Management Hosts tab. The Event Broker host status will be green.

Appendix A: Troubleshooting

This section includes material to help you troubleshoot problems or issues that may occur during the installation. Consult the Event Broker Administrator's Guide for additional detailed troubleshooting information.

Why do I see Failed to upload .. suite features ... failures when running uploadimages.sh during the installation?

You see this message "The suite-installer container is not running. Please make sure your suite-installer pod status is "RUNNING". Failed to upload the data of suite features."

Check that you are running the uploadimages.sh script from the correct folder, which is /opt/arcsight/kubernetes/scripts/.

What kind of errors can indicate potential DNS resolution issues?

DNS resolution issues can be indicated by the schema registry not running, and the schema registry pod in crash loop status, with following error message in the schema registry logs

```
# kubectl logs eb-schemaregistry-1138097507-1jxbn -n arcsighteventbroker
...
org.apache.kafka.common.config.ConfigException: No resolvable bootstrap urls
given in bootstrap.servers
...
```

How do I capture diagnostic data and logs?

Event Broker includes a range of diagnostic scripts in the web service container. For details on how to utilize these scripts, see Diagnostic Data and Scripts in the Event Broker Administrator's Guide.

Appendix B: Installer Command Line Arguments

The installer command line can include the following arguments.

```
./arcsight-installer-master.sh [--MASTER_NODES=<IPv4_Addr1> <IPv4_Addr2> <IPv4_Addr3>"] [--ROOTCA=<ca>] [--ROOTCAKEY=<cakey>] [--LOG_MAX_SIZE=<max_size_number>] [--LOG_MAX_FILE=<max_file_number>] [--DOCKER_HTTP_PROXY=<http_proxy>] [--DOCKER_HTTPS_PROXY=<https_proxy>] [--DOCKER_NO_PROXY=<no_proxy>] [--POD_CIDR=<pod_ip_range>] [--SERVICE_CIDR=<service_ip_range>] [--DNS_SVC_IP=<internal_dns_ip>] [--HA_VIRTUAL_IP=<virtual_ip>] [--EXTERNAL_ACCESS_HOST=<external_access_host>] [--NFS_SERVER=<nfs_ip_addr>|<internal>] [--NFS_FOLDER_ROOT=<root_nfs_folder>]
```

Argument	Description	Use with Cluster Type
--MASTER_NODES	The value is a list of the three master node IPv4 addresses, separated by a blank. The entire list is enclosed in one set of double-quotes. Three master nodes are recommended for production environments.	Multi-master
--REGISTRY_ORGNAME	The organization name in the local Docker registry where application images are located. The default value is 'arcsightsecurity'.	N/A This option does not need to be specified when running the script.
--ROOTCA	Specify the root CA certificate for generating server and client certificates.	Multi-master and Single master
--ROOTCAKEY	Specify the root CA key for generating server and client certificates.	Multi-master and Single master
--LOG_MAX_SIZE	The max file size used by Docker log rotation, default if not provided will be 20 MB.	Multi-master and Single master
--LOG_MAX_FILE	The max number of files used by the by Docker log rotation, default if not provided will be 5 files.	Multi-master and Single master
--DOCKER_HTTP_PROXY	Proxy settings for Docker. Specify if accessing the Docker hub or Docker registry requires a proxy. By default will be configured from http_proxy environment variable on your system.	Multi-master and Single master

Argument	Description	Use with Cluster Type
--DOCKER_HTTPS_PROXY	Proxy settings for Docker. Specify if accessing the Docker hub or Docker registry requires a proxy. By default will be configured from https_proxy environment variable on your system.	Multi-master and Single master
--DOCKER_NO_PROXY	Proxy settings for Docker. Specify if accessing the Docker hub or Docker registry requires a proxy. By default will be configured from no_proxy environment variable on your system.	Multi-master and Single master
--POD_CIDR	Kubernetes pod IP range. Default is 172.16.0.0/16.	Multi-master and Single master
--SERVICE_CIDR	Kubernetes service IP range. Default is 172.30.78.0/24. Must not overlap POD_CIDR range.	Multi-master and Single master
--DNS_SVC_IP	Kubernetes internal DNS service IP address. Must be within SERVICE_CIDR range. Default is 172.30.78.78	Multi-master and Single master
--HA_VIRTUAL_IP	A Virtual IP (VIP) is an IP address that is shared by all master machines. The VIP is used for the connection redundancy by providing fail-over for one machine. When a master goes down, the other master takes over the VIP address and responds to requests sent to the VIP. Mandatory for multi-master cluster.	Multi-master
--EXTERNAL_ACCESS_HOST	Defines a fully-qualified domain name of the Virtual IP address This is required for multi-master cluster deployments.	Multi-master
--NFS_SERVER	Persisted platform and product data will be stored in this location. The value is either the external NFS server IP, external NFS server hostname, or the text 'internal'. If the value is 'internal', a simple NFS server on master node will be installed. The internal option is supported only with single master cluster deployments. Multi-master deployments must use an external NFS server.	Multi-master and Single master
--NFS_FOLDER_ROOT	The root folder on the external NFS server, for example: /opt/arcsight/nfs/volumes The default 'internal' NFS location is '/opt/arcsight/volumes'	Multi-master and Single master
-h --help	Show help.	N/A

Appendix C: The arcsight-installer.properties file

The arcsight-installer.properties file controls several important settings for your Event Broker installation. You will need to adjust the default configuration setting if you are deploying in FIPS mode, or adding more worker nodes to the default configuration, etc. The settings are described here.

Before deployment,

- Edit the properties file, `/opt/arcsight/installer/arcsight-installer.properties`, as needed for your environment. To change optional property values, remove the comment operator (`#`) and then make the desired change.
- After changing values in the file, run the script `/opt/arcsight/installer/update-arcsight-installer-properties.sh` for the changes to take effect.

Setting	Notes
## All Event Broker components will use FIPS-certified encryption algorithms	
eb-init-fips=false	Turns FIPS on. Not recommended to change after deployment.
## Event Broker Kafka will use TLS Client Authentication to verify client connections	
eb-init-client-auth=false	Turns TLS-CA on. Not recommended to change after deployment.
## Number of partitions for Event Broker default topics in Kafka	
eb-init-noOfTopicPartitions=6	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	
eb-init-topicReplicationFactor=2	Default value. Will only affect newly created topics. (Must delete old topics to change replication factor.)
## Kafka log retention size	

Setting	Notes
eb-init-kafkaRetentionBytes=10737418240	Default value per partition per topic. Deletion will occur when Event Broker hits either the duration or retention bytes, whichever comes first.
## Kafka log retention size for the Vertica Avro topic. This is uncompressed and requires more space to hold events for the same duration.	
eb-init-kafkaRetentionBytesForVertica=10737418240	Default value per partition per topic. 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 Event Broker hits either the duration or retention bytes, whichever comes first.
## Kafka log retention duration	
eb-init-kafkaRetentionHours=672	Based on environment. Requires calculation on customer behalf. Applies to all topics, including those created through ArcMC. Deletion will occur when Event Broker hits either the duration or retention bytes, whichever comes first.
## The replication factor for the offsets topic	
eb-init-kafkaOffsetsTopicReplicationFactor=3	Defines the replication factor for the __consumer_offsets topic.
## The max time that Kafka waits to establish a connection to zookeeper	
eb-init-kafkaZKConnectionTimeoutMs=6000	Set the number of milliseconds that Kafka waits while attempting to establish a connection to Zookeeper.
## Zookeeper session timeout configuration for the Kafka broker	
eb-init-kafkaZKSessionTimeoutMs=6000	Set the number of milliseconds that zookeeper waits to receive a heartbeat from a broker or any consumer. If a heartbeat is not received in this period of time, the broker/consumer is assumed to be dead. A rebalance operation will be performed.

Setting	Notes
## Kafka inter-broker protocol version	
inter-broker-protocol-version=0.11.0.0	Only to be used during upgrades.
## The message format version the broker will use to append messages to the logs.	
log-message-format-version=0.11.0.0	Only to be used during upgrades.
## Number of Kafka and ZooKeeper nodes	
eb-kafka-count=3	Determines cluster size for Kafka. Must match number of worker nodes labeled as kafka=yes in Kubernetes. 1 node to 1 host.
eb-zookeeper-count=3	Determines zookeeper cluster size. Max of 7. Must match number of worker nodes labeled as zk=yes in Kubernetes. MUST be an odd number.
## Host path to store data persistently	
eb-kafka-path=/opt/arcsight/k8s-hostpath-volume/eb/kafka	Used if you have configured 'internal' NFS server. It will be created if it does not exist.
eb-zookeeper-path=/opt/arcsight/k8s-hostpath-volume/eb/zookeeper	Used if you have configured 'internal' NFS server. It will be created if it does not exist.
## ArcMC hostname	
eb-arcmc-hosts=localhost:443	
## The endpoint identification algorithm to validate the server hostname using the server certificate.	
ssl-endpoint-identification-algorithm=https	Hostname verification for Kafka to Kafka connections.
## The number of stream threads	
stream-num-threads=6	Do not change unless performance issue.
## truncate fields in C2av	

Setting	Notes
c2av-field-truncate=false	<p>Used by the CEF to Avro transformation stream processor used in the Investigate data pipeline. Defines how to handle events when the value for an attribute is longer than the maximum size defined by the Investigate Vertica schema.</p> <p>If set to false, the event will be rejected. It is not loaded to the Investigate events table, in Vertica. Instead, an entry is loaded to the rejected_events table.</p> <p>If set to true, the value will be truncated so that it fits the field limit. The event will be loaded to the Investigate events table.</p>
## c2av config params	Configuration properties for the CEF to Avro transformation stream processor used in the Investigate data pipeline.
c2av-heartbeat-interval-ms=1000	
c2av-max-poll-interval-ms=3600000	
c2av-max-poll-records=100	
c2av-session-timeout-ms=180000	
c2av-request-timeout-ms=305000	
## Log level for Event Broker components	Used to change the logging levels of each Event Broker component.
kafka-log-level=info	
zookeeper-log-level=info	
schema-log-level=info	
webservice-log-level=info	
c2av-stream-processor-log-level=info	
routing-processor-log-level=info	
## Host path directory for ArcMC certificates	Deprecated.
arcmc-certs-path=/opt/arcsight/k8s-hostpath-volume/eb/arcmccerts	In Event Broker 2.20, this is now configured using ArcSight Installer, under the Event Broker Configuration link > ArcMC Monitoring tab.

Setting	Notes
## Host path directory for AutoPass license data file persistence	
eb-autopass-path=/opt/arcsight/k8s-hostpath-volume/eb/autopass	Full path on the initial master, where the permanent license is stored. See the section 'Install your license before the evaluation period ends' or more information.

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 computer, click the link above and an email window opens with the following information in the subject line:

Feedback on Deployment Guide (Event Broker 2.20)

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!