



Hewlett Packard
Enterprise

HPE Security ArcSight Investigate

Software Version: 1.10

Deployment Guide

July 20, 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 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://community.saas.hpe.com/t5/Discussions/Third-Party-Copyright-Notices-and-License-Terms/td-p/1589228>

Support

Contact Information

Phone	A list of phone numbers is available on the HPE Security 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://community.saas.hpe.com/t5/ArcSight/ct-p/arc sight

Revision History

Date	Description
05/15/2017	Initial release of this document.

Contents

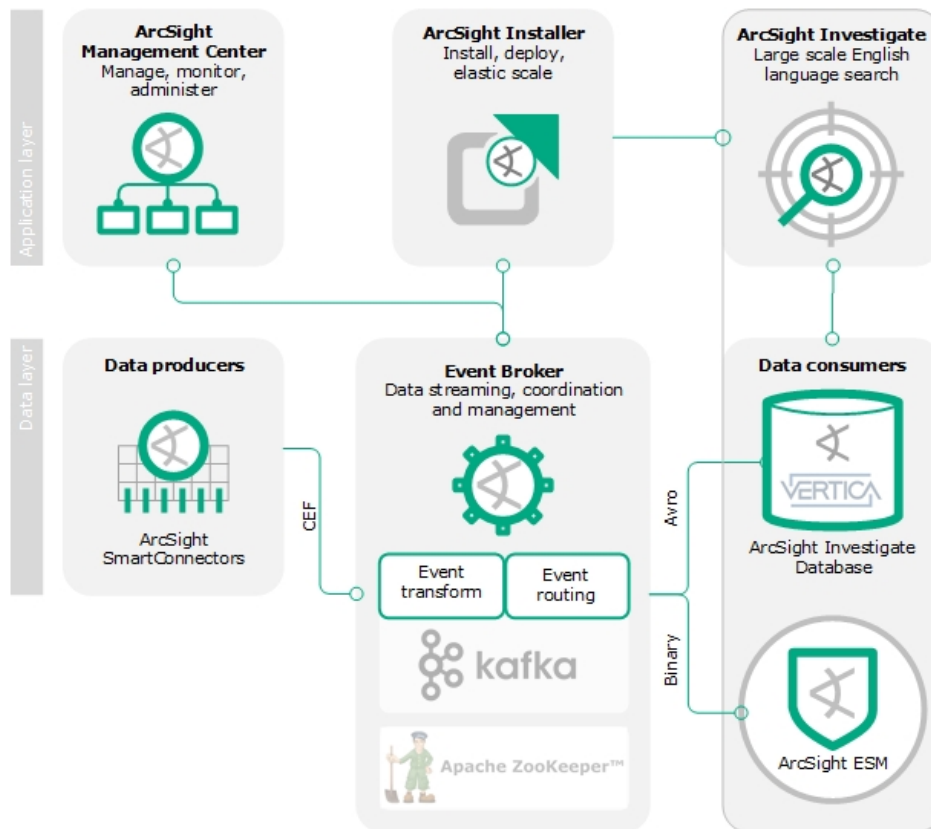
About ArcSight Investigate	6
About ArcSight Event Broker	7
Supported deployment scenarios	8
ArcSight Investigate deployment architecture	9
ArcSight Investigate deployment overview	11
What's next	11
ArcSight Investigate support matrix	11
Supported operating systems	11
Supported browsers	11
Supported product compatibility	12
What's next	12
ArcSight Investigate and Event Broker prerequisites	12
System requirements	12
Default heap size	14
Preparing producer and consumer interfaces and encryption modes	14
Set up encryption modes before installing and configuring Investigate and Event Broker	15
Preparing systems managed by Kubernetes and Vertica	16
TLS planning	17
Network planning	17
Setting security-enhanced Linux (SELinux) to permissive and disabling the firewall for Vertica	17
Setting up the proxy server	18
Installation order	20
What's next	20
Install ArcSight Investigate and Event Broker using the ArcSight Installer	20
Installing ArcSight Investigate offline (before deployment)	20
Generating a key pair on the master node for worker nodes	21
Installing ArcSight Installer on the master node	22

Adjust installer.properties before Kubernetes deployment	23
Deploying Kubernetes on the master node and setting up worker nodes	27
Deploying Investigate and Event Broker worker nodes in the ArcSight Installer	29
What's next	31
Generating a key pair on the Vertica cluster node 1 and installing the Investigate Vertica database ..	31
What's next	35
Configure ArcSight Investigate components	36
Configuring the ArcSight Investigate Vertica database connection in the ArcSight Installer	36
Configuring TLS on the ArcSight Investigate Vertica database in the Vertica server	36
Configuring Vertica SSL	37
Configuring the SMTP server in ArcSight Installer	38
Configure session and search settings in ArcSight Installer	38
Configure Event Broker for management by ArcMC	39
Configure SmartConnectors	39
Establishing the system admin	39
What's next	40
Generate signed certificates for consumers	40
Generate a signed certificate from the system CA	40
Generating a signed certificate from a CSR	41
Uninstalling ArcSight Investigate	41
ArcSight Investigate deployment troubleshooting and FAQs	42
Troubleshooting	42
Where to find the logs	42
Pod starting order	42
Cannot query zookeeper	42
Common Errors/Warnings in Zookeeper Logs	43
Common Errors/Warnings in Kafka logs	43
SSL Connection Error	43
The Event Broker pre-defined topics are not created on the initial Event Broker first deployment	44

One or more connectors cannot send data to Kafka	44
A consumer cannot read events from Kafka	45
An Event Broker component crashes: ArcMC Rest API, stream processors (Routing and Transform)	45
Event Broker EPS is lower than expected	45
FAQs	46
Which pods in Kubernetes comprise the Event Broker deployment?	46
Which pods in Kubernetes comprise the ArcSight Investigate deployment?	46
How can I find out more about Kafka and Apache Zookeeper?	46
Can I use my existing Event Broker v1.0 with ArcSight Investigate + Vertica?	47
With only a single install method for ArcSight Investigate and Event Broker, how can I distribute each of these (pods) across Kubernetes worker nodes?	47
How do I check the hostname of machines on which Event Broker is installed?	47
Appendix A: Adding Kubernetes worker nodes	48
Send Documentation Feedback	50

About ArcSight Investigate

ArcSight Investigate is a high-capacity data management and analysis engine that enables you to search, analyze, and visualize machine-generated data gathered from web sites, applications, sensors, and devices that comprise your monitored network. After you define the data source, Investigate indexes the data and parses it into events that you can view and search. You can use the English-like search language to generate results from which to create reports and visualizations.



Component	Description
ArcSight Investigate	High-capacity data management, search, and analysis web application.
ArcSight Installer	<p>A web application for deploying and configuring the ArcSight Investigate components, including ArcSight Investigate and Event Broker.</p> <p>The components are managed in a Kubernetes cluster. The master node hosts the ArcSight Installer web application and the ArcSight ArcSight Investigate web application, and the worker nodes host the Event Broker.</p>
ArcSight Investigate Vertica database	The ArcSight Investigate analytic database powered by Vertica provides high-capacity data storage and retrieval for rapid search response at high throughput. Vertica is installed separately.


Component	Description
ArcSight SmartConnectors	SmartConnectors collect and normalize event data from nodes on your network. Connectors normalize event 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. SmartConnectors can then filter and aggregate events to reduce the volume of events sent to the system. ArcSight SmartConnectors, installed and maintained separately, are producers that publish data to Event Broker. You can subscribe to data managed by Event Broker with ArcSight Investigate, ADP Logger, ArcSight ESM, Apache HDFS, or your own third-party consumer.
Event Broker	ArcSight Event Broker centralizes event processing, enabling you to take advantage of scalable, high-throughput, multi-broker clusters for publishing and subscribing to event data. Event Broker coordinates and manages data streams, which enables your ArcSight environment to scale, and opens up ArcSight events to third-party data solutions.
ArcMC	HPE ArcSight Management Center (ArcMC) is a centralized management tool that simplifies security policy configuration, deployment maintenance, and monitoring efficiently and cost-effectively. ArcMC provides run-time management of Event Broker topics. ArcMC is sold as part of the ArcSight Deployment Platform (ADP).

About ArcSight Event Broker

ArcSight Event Broker centralizes event processing and enables topic sorting and event routing, which helps you to scale your ArcSight environment, and opens 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 Apache Kafka. After you install and configure an Event Broker Kafka broker or cluster of broker nodes, you can use ADP SmartConnectors to publish data, and subscribe to that data with ADP Event Broker, ArcSight Investigate, ArcSight ESM, Apache Hadoop, or your own consumer.

Event Broker manages the distribution of events in topics to which consumers can subscribe.



- The CEF version configured at the connector on the producer side (CEF 0.1 or 1.0) should be the CEF version supported by the consumer.
 - CEF 0.1 is the legacy ArcSight CEF version that supports IPv4 IP addresses available with SmartConnector version 7.4 and earlier.
 - CEF 1.0, available with SmartConnector version 7.5 and later, supports IPv6 addresses.
- There are two Event Broker topics you can configure your SmartConnector to connect to: eb-cef and eb-esm. The ESM topic produces binary, which is the format ESM consumes.
- Multiple connectors can be configured to publish to the eb-cef topic. Load is balanced by Event Broker.
- Event Broker's stream processor converts CEF-formatted event data in the eb-cef topic to Avro format and sends the Avro formatted event data to eb-internal-avro topic.
- ArcSightESM can be configured as an Event Broker consumer, and can send query parameters to ArcSight Investigate via https.
 - The SmartConnector supporting this scenario has to be set up to dual-feed to both the eb-cef topic that connects to the ArcSight Investigate Vertica database and the eb-esm topic that connects to ESM.
 - This ensures that the exact same data generated from the connector is stored in both the Investigate database and the ESM database so that the queries sent by ESM to Investigate will produce consistent results in ArcSight Investigate.

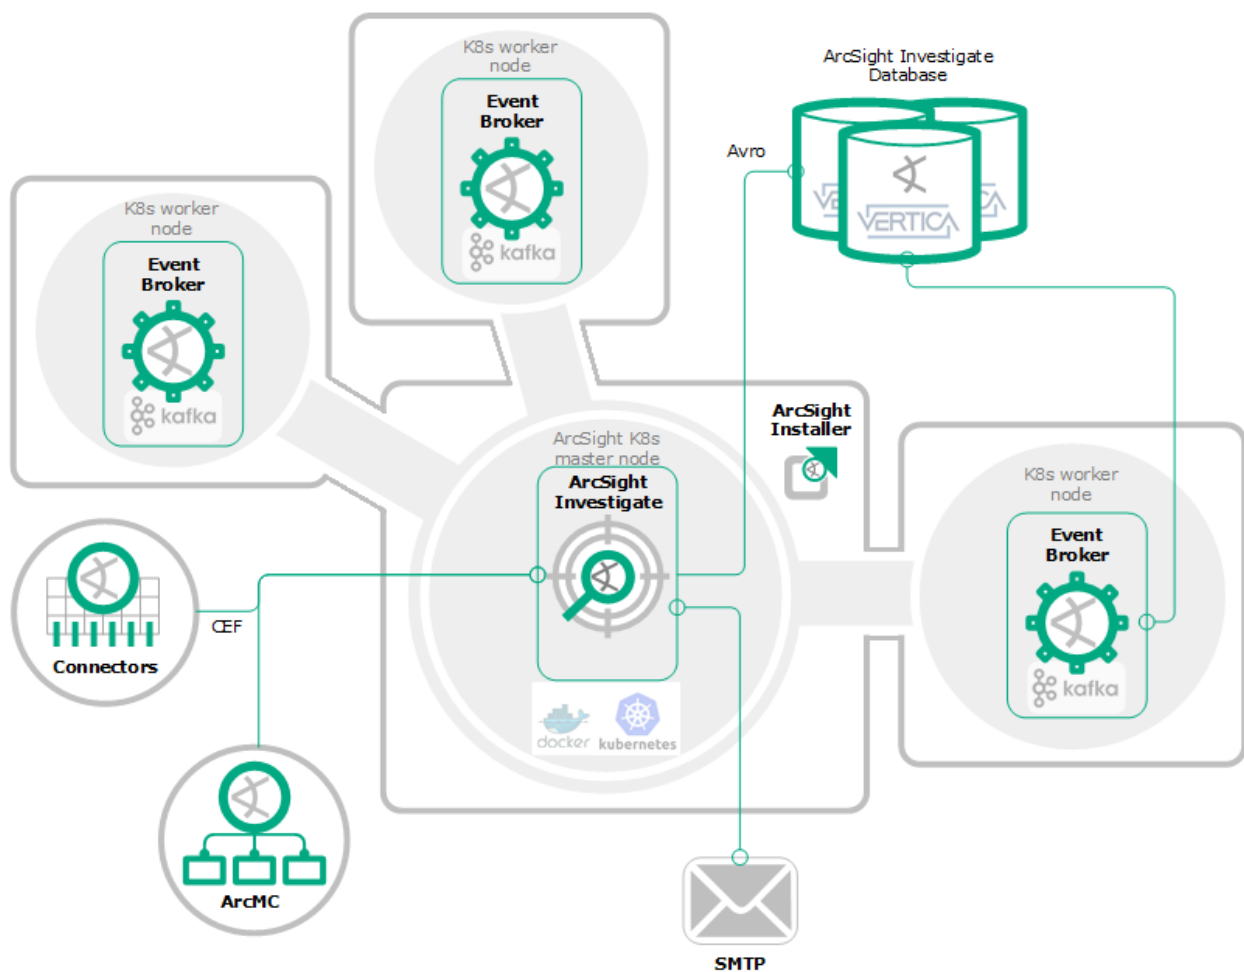
Supported deployment scenarios

ArcSight Investigate supports the following deployment scenarios:

	Use Case	Description	Guidance document
1	Fresh install	Install Investigate and Event Broker using ArcSight Installer <ul style="list-style-type: none"> standard encryption and FIPS encryption mode 	<i>Investigate Deployment Guide</i>
2	Existing ADP 2.0 customer with Event Broker 1.0	<ol style="list-style-type: none"> Migrate Event Broker 1.0 data to Event Broker 2.0 (check with Customer Support for the availability of data migration instructions) upgrade ArcMC to 2.6 Install Investigate from ArcSight Installer app 	<ol style="list-style-type: none"> <i>Event Broker data migration tech note</i> <i>ArcMC Administration Guide</i> <i>Investigate Deployment Guide</i>
3	Existing ADP 2.0 customer without Event Broker 1.0	<ol style="list-style-type: none"> upgrade to ArcMC to 2.6 install Investigate from ArcSight Installer app 	<ol style="list-style-type: none"> <i>Investigate Deployment Guide</i> <i>ArcMC Administration Guide</i>
4	Stand-alone Event Broker 2.0 without Investigate	Fresh install using stand-alone Event Broker version of ArcSight Installer app	<i>Event Broker Administration Guide</i>
5	Stand-alone Event Broker 1.0 to Event Broker 2.0 data migration	Instructions for migrating data from Event Broker 1.0 to Event Broker 2.0	<i>Event Broker data migration tech note</i>

ArcSight Investigate deployment architecture

Investigate is installed using Docker containers managed by Kubernetes and deployed from the ArcSight Installer application. The default deployment consists of a Kubernetes master node, and three Kubernetes (K8) worker nodes: three worker nodes for ArcSight Event Broker, and one Kubernetes master node for ArcSight Investigate.



Deployment component	Host	Functional contents
Kubernetes master node	1 VM or physical server	<ul style="list-style-type: none"> Kubernetes master node ArcSight Investigate ArcSight Installer application
Kubernetes worker nodes	3 VMs or physical servers	<ul style="list-style-type: none"> 3 Event Broker nodes
Vertica database	3 physical servers	<ul style="list-style-type: none"> 3 ArcSight Investigate Vertica database instances
ArcSightSmartConnectors	Stand-alone or part of ArcMC	Normalizes event data from network devices and formats as CEF.
ArcSight Management Console	Separate installation	Provides run-time management of Event Broker topics.
SMTP server	Separate installation	Provides the ability for ArcSight Investigate to send notification messages to users.

ArcSight Investigate deployment overview

1. ArcSight Investigate and Event Broker prerequisites
2. Install Investigate Vertica database
3. Install Investigate and Event Broker using the ArcSight Installer application
4. Configure Investigate components
5. Adding Kubernetes nodes
6. Uninstalling Investigate
7. Investigate deployment troubleshooting and FAQs

What's next

[ArcSight Investigate and Event Broker prerequisites](#)

ArcSight Investigate support matrix

Supported operating systems

Version	Component	Operating System
1.10	Investigate	RHEL 7.3 64-bit CentOS 7.3 64-bit
	Kubernetes cluster servers	Linux kernel version 3.10 or higher
	ArcSight Investigate Vertica 8.0.1-5 database	RHEL 7.0 and CentOS 7.0

Supported browsers

Browser	Version
Microsoft Edge	latest
Microsoft Internet Explorer	latest
Google Chrome	latest
Mozilla Firefox	latest

Supported product compatibility

Product	Version
ArcSightSmartConnector	7.5 and later
ArcMC	2.6
ArcSightLogger	6.4
ArcSightESM	6.11

What's next

[ArcSight Investigate and Event Broker prerequisites](#)

ArcSight Investigate and Event Broker prerequisites

System requirements

This information provides general sizing guidelines based on a default setup. For tailored sizing recommendations for a production environment, contact ArcSight Customer Support.

Provision the servers (or VMs) that you are using for the deployment with the following. For supported platforms and operating systems, see the ArcSight Investigate Support Matrix.

Component	Nodes	Resources needed	Needed ports
ArcSight Investigate + Event Broker	1 master 3 worker nodes	<ul style="list-style-type: none"> One CPU with 24 cores 32 GB RAM 8 TB disk space Linux kernel version 3.10 or higher Java (OpenJDK) 1.8.0_121 or higher Method for obtaining Docker containers, either via Internet (or proxy) or other internal method 10 GigE network 	<p>ArcSight Installer: 8888</p> <p>Kubernetes: 2379, 2380, 4001, 4194, 5000, 5443, 8080, 8088, 8200, 8285, 8443, 10248-10252, 10255, 30001</p> <p>Network File System (NFS): 111, 2049, 20048, 37189</p> <p>For required Event Broker ports, see "System requirements" in the <i>HPE Security ArcSight Data Platform Event Broker Deployment Guide</i>.</p>
ArcSight Investigate Vertica database	3	<p>Important: The Vertica database must be installed on the same sub-network as the ArcSight Investigate master and worker nodes.</p> <ul style="list-style-type: none"> 2 CPUs with 24 cores 128 GB RAM 8TB disk space 10 GigE network minimum (dual recommended) <p>Recommendation: Install Vertica on a dedicated physical server, for example HPE Proliant G9 or similar</p> <p>Virtual environment: HPE Vertica performs better on a physical server than in a virtualized environment because of the overhead and resource constraints imposed by the virtualization software. See HPE Vertica Analytics Platform Version 8.0.x Documentation for more information.</p>	5433

Component	Nodes	Resources needed	Needed ports
ArcMC (part of ADP)	1	<ul style="list-style-type: none">• One CPU quad-core• 16 GB RAM• 50 GB of free disk space <p>For ArcMC deployment details, see the <i>ArcMC Administrator's Guide</i>.</p>	
SmartConnectors (part of ADP)	1	<p>SmartConnector version 7.5 (can be stand-alone or managed by ArcMC)</p> <p>For ArcSightSmartConnector deployment details, see the <i>SmartConnector User's Guide</i>.</p>	

Default heap size

Following are the heap memory usage settings for Event Broker modules at the JDK level. These levels are not configurable.

Module Name	Default heap sizes
Schema Registry	1 GB
Kafka	4 GB
Kafka Manager	1 GB
c2a SP	2 GB
Routing SP	2 GB
Web Service	2 GB

Preparing producer and consumer interfaces and encryption modes

Follow these guidelines for preparing optional producer and consumer components. See the indicated guidance documentation for detailed instructions. For instructions about how to configure producers and consumers after ArcSight Investigate and Event Broker deployment, see [Configure ArcSight Investigate components](#).

Set up encryption modes before installing and configuring Investigate and Event Broker

Before installing Investigate and Event Broker, determine the encryption mode you want to use to encrypt communications between ArcSight components. Set up the other ArcSight components with the encryption mode you intend to use before connecting them to the Event Broker. The security mode of systems connected to Event Broker (Consumers, Producers, ArcMC) must be the same as the security mode set for Event Broker. Changing encryption modes 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 encryption modes	Guidance documentation
ArcMC	Install ArcMC before ArcSight Investigate and Event Broker installation.	38080	<ul style="list-style-type: none">• TLS• FIPS• ClientAuth	<i>ArcMC Administrator's Guide</i>
ArcSightSmartConnectors	<p>ArcSightSmartConnectors and ArcMC onboard connectors can be installed and running prior to installing ArcSight Investigate and Event Broker.</p> <p>FIPS mode setup is not supported between Connector version 7.5 and Event Broker. TLS and ClientAuth are the only encryption methods supported between SmartConnector version 7.5 and Event Broker.</p>	39093	<ul style="list-style-type: none">• TLS• ClientAuth• FIPS	<i>SmartConnector User Guide</i> <i>ArcMC Administrator's Guide</i>

Product	Preparations needed	Open ports	Supported encryption modes	Guidance documentation
ArcSightESM (optional)	ArcSightESM can be installed and running prior to installing ArcSight Investigate and Event Broker. ESM ingests events faster than Investigate does. (Investigate Scheduler ingests events at 22K per second while ESM ingests events at 30K per second.) You can leave the ingestion rate asynchronous, or you can equalize them by setting the ESM ingestion rate to a lower rate at the connector so that Investigate and ESM ingest rates are closer. This will reduce the likelihood of a lag in search results on Investigate launched from ESM.	39093	<ul style="list-style-type: none">• TLS• FIPS• ClientAuth	<i>ESM Installation Guide</i> <i>ESM Administrator's Guide</i>
ArcSightLogger (optional)	ArcSightLogger can be installed and running prior to installing Event Broker.	39093	<ul style="list-style-type: none">• TLS• FIPS• ClientAuth	<i>Logger Administrator's Guide</i>

Preparing systems managed by Kubernetes and Vertica

Procedure

Prerequisite: Prepare the Kubernetes master and worker nodes ("[Preparing producer and consumer interfaces and encryption modes](#)" on page 14)

- Configure NTP using Chrony on all of the hosts in the cluster.
Chrony is installed by default on some versions of Red Hat and CentOS.

If you do not have Chrony installed, run the following commands:

- a. Install Chrony:

```
yum install chrony
```

- b. Start Chronyd:

```
systemctl start chronyd  
systemctl enable chronyd
```


- c. Verify that Chrony is operating correctly:

```
chronyc tracking
```

TLS planning

The various components in the ArcSight Investigate system interact using encrypted communication implemented using TLS 1.2 protocol.

TLS implementation requires digital certificates. Before you begin the installation process, you must decide on the type of certificate you prefer to use:

- Kubernetes self-signed certificate. Kubernetes includes the capability to generate self-signed certificates. By default, the Kubernetes installation process generates certificates for the Kubernetes cluster, but you can instruct otherwise during the installation process. You can also generate a Kubernetes certificate for other components in the system, which require a certificate, like the ArcSight Investigate Vertica database. For more information on generating a Kubernetes certificate, see [Generate signed certificates for consumers](#).
- A valid digital certificate signed by a certificate authority (CA). Depending on your organization's security policy, you might be required to use a certificate from a trusted CA. In this case, make sure that you have a root certificate file and a private key file. Copy these files to the designated Kubernetes master node.

Note: The certificates cannot be reconfigured after installation.

If you are planning on enabling FIPS mode, make sure the certificate generated meets the FIPS criteria.

Network planning

- Ensure that each node is configured with a fully qualified domain name
- Ensure proper DNS configuration across all systems including correct forward and reverse proxy lookups
- Enable internet access in order to download the product container images.
- If your organization's network has a proxy, define the proxy environment variable on all servers.

Setting security-enhanced Linux (SELinux) to permissive and disabling the firewall for Vertica

About

This procedure is needed for the Vertica host only.

Procedure

1. Set SELinux to permissive.
 - a. SELinux status is enabled by default. To check status, check the file `/etc/sysconfig/selinux`:


```
vi /etc/sysconfig/selinux <command>
```
 - b. If SELinux=enforcing, then change to SELinux=permissive.
 - c. Save and exit the file.
 - d. Reboot the server.
2. Disable firewall on the Vertica system.
 - a. To check firewall status, run the following command on the operating system as a root:


```
systemctl list-unit-files | grep firewall
```
 - b. If return status is “firewalld.service enabled”, then run the following command:


```
systemctl stop firewalld
```



```
systemctl disable firewalld
```

Setting up the proxy server

About

If you use a proxy server in your environment, define a proxy environment variable on all servers.

Procedure

1. Configure proxy settings (if needed) to get yum access to CentOS repos.


```
cat >> ~/.bashrc <<EOL export ftp_proxy=<your_proxy_url:port> export http_
proxy=<your_proxy_url:port> export https_proxy=<your_proxy_url:port> EOL
source ~/.bashrc
```
2. Increase the default user process.
 - a. Open the file `/etc/security/limits.d/20-nproc.conf`.
 - i. If you do not already have a `/etc/security/limits.d/20-nproc.conf` file, create one (and the `limits.d` directory, if necessary).
 - ii. If the file already exists, update and add entries in the file.
 - b. Add the lines below.

Be sure to include the asterisk * in the new entries. It is important that you add all of the entries exactly as specified.

```
* soft nproc 10240
* hard nproc 10240
* soft nofile 65536
* hard nofile 65536
```

3. Reboot all Kubernetes master and worker nodes, and Vertica nodes.

Nodes can be rebooted in any order.

4. Verify that all nodes are up and running.

Example:

```
#
[root@n11-222-333-h111 ~]# ulimit -a
core file size          (blocks, -c) 0
data seg size           (kbytes, -d) unlimited
scheduling priority     (-e) 0
file size               (blocks, -f) unlimited
pending signals         (-i) 514466
max locked memory       (kbytes, -l) 64
max memory size         (kbytes, -m) unlimited
open files              (-n) 65536
pipe size               (512 bytes, -p) 8
POSIX message queues    (bytes, -q) 819200
real-time priority      (-r) 0
stack size              (kbytes, -s) 8192
cpu time                (seconds, -t) unlimited
max user processes      (-u) 10240
virtual memory          (kbytes, -v) unlimited
file locks              (-x) unlimited
[root@n11-222-333-h111 ~]# sestatus
SELinux status:          disabled
[root@n11-222-333-h111 ~]# systemctl status firewalld
● firewalld.service - firewalld - dynamic firewall daemon
   Loaded: loaded (/usr/lib/systemd/system/firewalld.service;
   enabled; vendor preset: enabled)
   Active: active (running) since Fri 2017-03-17 16:21:52 PDT; 7min ago
     Docs: man:firewalld(1)
   Main PID: 1223 (firewalld)
    CGroup: /system.slice/firewalld.service
           └─1223 /usr/bin/python -Es /usr/sbin/firewalld --nofork --nopid
Mar 17 16:21:51 n11-222-333-h111.domainname.com systemd[1]: Starting
firewalld - dynamic firewall daemon...
Mar 17 16:21:52 n11-222-333-h111.domainname.com systemd[1]: Started
firewalld - dynamic firewall daemon.
```

```
[root@n11-222-333-h111 ~]# hostname -f  
n11-222-333-h111.domainname.com
```

Installation order

You can install ArcMC and ArcSightSmartConnectors in any order during ArcSight Investigate deployment, however, ArcSight recommends having them installed and running in your environment before installing the ArcSight Investigate components. Additional configuration of both components will be required after you install ArcSight Investigate.

1. Install ArcSight Investigate and Event Broker using the ArcSight Installer application
2. Install ArcSight Investigate Vertica database
3. Configure ArcSight Investigate components

What's next

Generating a key pair on the Vertica cluster node 1 and installing the Investigate Vertica database

Install ArcSight Investigate and Event Broker using the ArcSight Installer

Before running the ArcSight Installer application, verify that you have set up the receiving systems according to guidelines in ArcSight Investigate and Event Broker prerequisites. The ArcSight Installer will configure firewall settings during setup (in case `firewalld.service` is up and running) on both the Kubernetes master and worker nodes. Multi-master installation is not supported.

This procedure provides instructions for installing online using the Docker Hub repository, or offline by downloading a tar file from an FTP site and replicating a local Docker Hub on the master node system.

Installing ArcSight Investigate offline (before deployment)

About

Installing ArcSight Investigate offline requires no internet connection from your Investigate servers to the HPE Docker Hub registry. You can download the Docker files from the HPE Software Depot:

<https://h20392.www2.hpe.com/portal/swdepot/index.do>.

Installing offline requires the additional `arcsight-eventbroker-2.02.0-images_66_4f7e26.tar` and `arcsight-investigate-1.10.0-images_66_4f7e26.tar` files.

Procedure

1. From a server with an internet connection, connect to <https://h20392.www2.hpe.com/portal/swdepot/index.do> and download the installer files:

Component	Offline installer file name
ArcSight Installer	arcsight-installer-1.10.7-ga110.x86_64.rpm Note: File is for both online and offline installations.
ArcSight Event Broker offline images	arcsight-eventbroker-2.02.0-images_66_4f7e26.tar arcsight-investigate-1.10.0-images_66_4f7e26.tar Note: File is only for offline installation.
ArcSight Investigate offline images	arcsight-investigate-1.10.0-images_66_4f7e26.tar Note: File is only for offline installation.
ArcSight Investigate Vertica Database	arcsight-investigate-vertica-scripts.136ef34.tar.gz

2. Copy the rpm and tar files to any location on the master node.
3. Push the images to the private local registry.

A private Docker registry is configured and running on 127.0.0.1:5000 and is accessible from all nodes and masters.

Run the following command on the master node for the Investigate image and Event Broker image:

```
/opt/arcsight/k8s/master/pushImages.sh -f <images.tar>
```

Where `<images.tar>` is the following:

- `arcsight-eventbroker-2.02.0-images_66_4f7e26.tar`
- `arcsight-investigate-1.10.0-images_66_4f7e26.tar`

Generating a key pair on the master node for worker nodes

About

In a master and worker node deployment, generate a key pair on the master node and then copy the public key to each worker node. This enables password-less SSH access from the master server to all the other worker node servers in the cluster. Do this before you install the ArcSight Installer, and before you install and setup Kubernetes.

The following is an example of enabling password-less SSH. For additional examples, see http://www.linuxproblem.org/art_9.html

Note: Generate the key pair as a root user or pseudo user.

Procedure

1. Run the `ssh-keygen` command on the master server.

Example:

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

2. Copy the key from the master node to the worker node using the worker node's IP address.

Example:

```
ssh-copy-id -i ~/.ssh/id_rsa.pub root@11.111.111.111
```

The system displays the key fingerprint and requests to authenticate with the node server.

3. Enter the worker node credentials as required.

The operation is successful when the system displays the following message:

```
Number of key(s) added: 1
```

4. To verify that the key was successfully installed on the worker node, run the following command from master to the worker node to verify that it can successfully log into the worker node.

```
ssh 'root@11.111.111.111'
```

5. Repeat steps 2 through 4 for every worker node.

Installing ArcSight Installer on the master node

About

Use this procedure to install the ArcSight Installer, starts the application, and unpacks the Kubernetes scripts to `/opt/arcsight/installer/k8s`.

Note: Perform the installation as a root user or pseudo user.

Procedure

1. Copy the installation package to the master server.

2. Verify that the tarball matches the md5 checksum:

```
md5sum arcsight-installer-1.10.7-ga110.x86_64.rpm
```

```
cat arcsight-installer-1.10.7-ga110.x86_64.rpm.MD5
```

Both outputs should match.

3. Run the following command on the master server, for example, `arcsight-installer-1.10.7-ga110.x86_64.rpm`

```
yum install -y arcsight-installer-1.10.7-ga110.x86_64.rpm
```

4. To verify that the ArcSight Installer was installed successfully, open the ArcSight Installer URL in a browser on the master server.

`https://<kubernetes_master>:8888/`

The ArcSight Installer login screen should load in the browser. If you are using self-signed certificates, accept the self-signed certificate notification from the browser. If the ArcSight Installer login screen does not load, see ArcSight Investigate deployment troubleshooting and FAQs.

The Add New Cluster dialog appears. Before you can specify information for this dialog, complete ["Deploying Kubernetes on the master node and setting up worker nodes" on page 27](#).

Adjust installer.properties before Kubernetes deployment

Before deploying Kubernetes, adjust properties, `/opt/arcsight/installer/installer.properties`, as needed for your environment. You would need to adjust the properties set here if you are deploying in FIPS mode, or want to add more worker nodes to the default configuration.

Property file setting	Default value	Options
All Event Broker components will use FIPS-certified encryption algorithms	<code>predeploy.eb.init.fips=false</code>	Set to true if you are setting up FIPS encryption
Event Broker Kafka will use TLS Client Authentication to verify client connections	<code>predeploy.eb.init.client-auth=false</code>	Set to true if you are setting up client-auth authentication
Number of partitions for Event Broker topics in Kafka	<code>predeploy.eb.init.noOfTopicPartitions=6</code> where: $\text{partitions} = n * m$ $n = \text{any number}$ $m = \text{total number of Vertica servers (default is three)}$ The number of Vertica servers can be increased by increments of three, such as 6, 9, 12.	Set a different number of partitions as needed

Property file setting	Default value	Options
Replication factor for Event Broker topics in Kafka	predeploy.eb.init.topicReplicationFactor=2	Set a different replication factor for Kafka as needed. Example: Replication factor of 2 means that 2 copies of the data are kept across the cluster for redundancy. Replication factor of 1 means only 1 copy of the data. It is common to set the property to 2 or 1.
Kafka log topic partition retention size (bytes)	predeploy.eb.init.kafkaRetentionBytes=10737418240	Set a different size (in bytes) for the Kafka log topic partition retention size as needed. The default retention size for events is 10 GB. Increase this value if you want to keep more events in Event Broker.

Property file setting	Default value	Options
Kafka log topic partition retention size for the Vertica Avro topic. This is uncompressed and requires more space to hold events for the same duration. (bytes)	predeploy.eb.init.kafkaRetentionBytesForVertica=10737418240	Set a different size (in bytes) for the Vertica Avro topic partition as needed. The default retention size for events is 10 GB. Increase this value if you want to keep more events in Event Broker.
Kafka log retention duration (hours)	predeploy.eb.init.kafkaRetentionHours=672	Set a different duration for the Kafka log retention as needed. By default, events are kept on Event Broker for 28 days, if the retention space is not reached first. You can raise or lower the retention period.
Kafka inter-broker protocol version	predeploy.inter.broker.protocol.version=0.10.1.0	
The message format version the broker will use to append messages to the logs.	predeploy.log.message.format.version=0.10.1.0	
Size of kafka and zookeeper pet-sets	predeploy.eb.kafka.count=3 predeploy.eb.zookeeper.count=3	
Host path to store data persistently	predeploy.eb.kafka.path=/opt/arcsight/k8s-hostpath-volume/eb/kafka predeploy.eb.zookeeper.path=/opt/arcsight/k8s-hostpath-volume/eb/zookeeper	

Property file setting	Default value	Options
ArcMC hostname	predeploy.eb.arcmc.hosts=localhost:443	This setting is populated during setup in ArcMC.
The endpoint identification algorithm to validate the server hostname using the server certificate.	predeploy.ssl.endpoint.identification.algorithm=https	If hostname validation is not available, no value is needed.
The number of stream threads	predeploy.stream.num.threads=6	
truncate fields in C2av	predeploy.c2av.field.truncate=false	Change to true to avoid events being sent to Vertica's reject table in case the event field's information size exceeded the defined length.
Log level for each Event Broker container	predeploy.level=info predeploy.kafka.log.level=\${predeploy.level} predeploy.zookeeper.log.level=\${predeploy.level} predeploy.schema.log.level=\${predeploy.level} predeploy.web.service.log.level=\${predeploy.level} predeploy.c2av.stream.processor.log.level=\${predeploy.level} predeploy.eventbroker.routing.processor.log.level=\${predeploy.level}	For debugging, change "info" to "debug"
Host path directory for ArcMC certificates	predeploy.arcmc.certs.path=/opt/arc sight/ k8s-hostpath-volume/eb/arcmc certs	This setting is populated during Event Broker setup in ArcMC

Deploying Kubernetes on the master node and setting up worker nodes

Procedure

Prerequisite: Update `/opt/arcsight/installer/installer.properties` ("Adjust [installer.properties before Kubernetes deployment](#)" on page 23)

1. Perform the installation as a root user.
2. Run the following command on the master server:

```
/opt/arcsight/installer/k8s/master/install.sh [optional parameters]
```

Optional parameters: If you are using a certificate from a trusted CA, pass the following parameters:

ROOTCA: The root or intermediate certificate for generating client and server certificates.

ROOTCAKEY: The private key for generating the client and server certificates.

CLOUD_PROVIDER: This parameter is required if your cloud provider is Microsoft Azure, which requires a specific configuration. Pass the following value: azure.

Example for ROOTCA and ROOTCAKEY:

```
/opt/arcsight/installer/k8s/master/install.sh ROOTCA=/tmp/ca.crt  
ROOTCAKEY=/tmp/ca.key
```

Example for CLOUD_PROVIDER

```
/opt/arcsight/installer/k8s/master/install.sh CLOUD_PROVIDER=azure
```

3. Override the default configuration of the Docker log rotation (Docker log rotation keeps a total of 5 files per container, each file with maximum of 20 MB):

```
/opt/arcsight/installer/k8s/master/install.sh LOG_MAX_SIZE=100m LOG_MAX_FILE=5
```

LOG_MAX_SIZE: specifies the maximum log file size, per container, for Docker log rotation, indicating a number followed by the unit of measure (k=Kilobytes, m=Megabytes, and g=Gigabytes).

LOG_MAX_FILE: specifies the maximum number of files to keep, per container, for Docker log rotation.

4. Answer the following prompts.

Prompt-1:

If desired IPv4 address is not displayed below use Ctrl+C to exit and make sure corresponding network interface is up before running this script again.

1) 15.214.136.220

#? 1

Prompt-2:

Note: For an offline installation, press **Enter** to dismiss the prompt. Responding to this prompt will cause the offline installation to fail.

If proxy is used, please define it in form `http://proxy.example.com:80/`
`http://web-proxy.corp.hpecorp.net:8080/`

5. Run the following command on the master server for each of the worker nodes. Worker node IP should be specified in a IPv4 format (1.1.1.1).

```
/opt/arcsight/installer/k8s/node/install.sh -w <worker node IP>
```

6. Label every Kubernetes node so that Search will always be assigned to the Master node, and so that data of each Kafka maps to its host server to avoid data loss.

The default Event Broker configuration assumes a cluster has three workers. This also means that every worker node will host an instance of Kafka and Zookeeper. So there will be three instances of Kafka and three instances of Zookeeper running in the cluster. Every Kubernetes node should be manually labeled using the command (repeat this procedure on every Kubernetes node):

For the default 4-server setup:

```
Master/Search: kubectl label --overwrite node <master-node-ip>
investigate=yes
```

```
Worker1/kafka/zookeeper: kubectl label --overwrite node <work1-node-ip>
zk=yes kafka=yes
```

```
Worker2/kafka/zookeeper: kubectl label --overwrite node <work2-node-ip>
zk=yes kafka=yes
```

```
Worker3/kafka/zookeeper: kubectl label --overwrite node <work3-node-ip>
zk=yes kafka=yes
```

7. Create a new cluster.
 - a. Open the ArcSight Installer on the master server:


```
https://<kubernetes_master>:8888
```
 - b. On the Login page click **Create New Cluster**.
 - c. On the Add New Cluster page, enter the following information and then click **Create**:
 - **Cluster ID.** Enter an ID made up of letters or letters and numbers or underscore character, and must be at least 5 characters in length. Spaces are not supported.
 - **Password/Confirm Password.** Enter a password that is at least six characters long, and then confirm it.
8. Set up the connections to the master Kubernetes node and Docker Hub.
 - a. From the Cluster Setup page of the ArcSight Installer, click **Edit** next to the **Master** field.
 - b. On the Master Configuration page, enter the following information and then click **Save**:

- **MasterAddress.** Enter the hostname or IP address of the master server.
- **Token.** Copy the content of the token into this field from the following location in the master node:
`/opt/arcsight/kubernetes/ssl/token`

Note: Remove the `,admin,admin` portion of the location path.

- **CA Certificate.** Copy the content of the CA certificate into this field from the following location in the master node:
`/opt/arcsight/kubernetes/ssl/ca.crt`

9. Configure the Docker hub.

On the **Cluster setup page** > **Docker Repository** field, click **Edit**.

- If you are installing offline, enter `127.0.0.1:5000` in the **URL** field and enter dummy data in the **UserName**, **Password**, and **email** fields (see "[Installing ArcSight Investigate offline \(before deployment\)](#)" on page 20).
- If you are installing online, specify real values in the dialog for Docker hub.

Connect to `index.docker.io` and access your Docker Investigate account.

10. From the ArcSight Installer, click **Node Management** and verify that the servers that you installed have **READY** status.

Note: Allow several minutes for the established nodes to appear as **READY** on the Node Management page.

If the nodes are not ready after several minutes, refresh the page. If the ready nodes still do not appear, see [ArcSight Investigate deployment troubleshooting and FAQs](#).

Status	Description
NOT_READY	Node is available but not ready for product deployment.
READY	Node is available and ready for product deployment.
ERROR	Node is available, but cannot be deployed due to a system error. Hover over the error icon to see the message.

Deploying Investigate and Event Broker worker nodes in the ArcSight Installer

Procedure

1. In the ArcSight Installer, click the **Deployment** tab.
2. Click **Deploy** next to Event Broker.

You can monitor the deployment status in the status column. The initial status is **IN_PROGRESS**. Deployment may take a while depending on the connection speed to the Docker Hub repository. When the products are deployed, the status changes to **DEPLOYED**.

Status	Description
NOT_READY	ArcSight Installer cannot communicate with Kubernetes. Deployment is not possible.
OFF	ArcSight Installer is able to communicate with Kubernetes. The products are not deployed yet.
IN_PROGRESS	Deployment or un-deployment has started and is in progress. This status can also display when one or more containers get restarted on Kubernetes (for example, when you change product configurations).
DEPLOYED	Product is successfully deployed and running.
ERROR	One or more product containers is broken. Hover over the error icon to see the message. This status may show up and then turn into DEPLOYED when a container has crashed and then fixed or restarted by Kubernetes.

- Click **Deploy** next to ArcSight Investigate.

You can monitor the deployment status in the status column. The initial status is **IN_PROGRESS**. Deployment may take a while depending on the connection speed to the Docker Hub repository. When the products are deployed, the status changes to **DEPLOYED**.

Status	Description
NOT_READY	ArcSight Installer cannot communicate with Kubernetes. Deployment is not possible.
OFF	ArcSight Installer is able to communicate with Kubernetes. The products are not deployed yet.
IN_PROGRESS	deployment or undeployment has started and is in progress. This status can also display when one or more containers get restarted on Kubernetes (for example, when you change product configurations).
DEPLOYED	the product is successfully deployed and running.
ERROR	one or more product containers is broken. Hover over the error icon to see the message. This status may show up and then turn into DEPLOYED when a container has crashed and then fixed or restarted by Kubernetes.

- Verify the integrity of the files using the Docker inspect command. Investigate includes 6 Event Broker images and 4 Investigate images. Event Broker includes 6 Event Broker images.

```
$ docker inspect --format='{{.Id}}' <images>
```

- Example output for the 6 Event Broker images:

```
$ docker inspect --format='{{.Id}}' arcsightsecurity/atlas_
kafkamanager:2.02.0
```

```
sha256:b1849f64a427e64c896b9070b708a48f053bda4828b04638bab2adce41cc0f48
```

```
$ docker inspect --format='{{.Id}}' arcsightsecurity/atlas_web-service:
```

2.02.0

```
sha256:b0e0b734ed54ca02f49f927a1e0af6e43423f8f5a91effff03ecaac11ed8a57b
```

```
$ docker inspect --format='{{.Id}}' arcsightsecurity/atlas_sp:2.02.0
```

```
sha256:5d51459969571501a798d9d7ba798b579da44fd9aa98d5892a4ac3505f444ca9
```

```
$ docker inspect --format='{{.Id}}' arcsightsecurity/atlas_
schemaregistry:2.02.0
```

```
sha256:d79b406bda0573489d185bb587db567d564a18047204c5474638dbb0066c6583
```

```
$ docker inspect --format='{{.Id}}' arcsightsecurity/atlas_kafka:2.02.0
```

```
sha256:4e8bcf9899e02c741cae1aa21825d77da3bf3c29061b49bc250bb29d7bb9ea9d
```

```
$ docker inspect --format='{{.Id}}' arcsightsecurity/atlas_
zookeeper:2.02.0
```

```
sha256:cb96c4e0e960e26b2ba0b66e8ed195df9d1a1024fbdb312e52d23c9b751da527
```

b. Example output of the 4 Investigate images:

```
$ docker inspect --format='{{.Id}}' arcsightsecurity/search:1.10.0
```

```
sha256:05341c212fab6bafc343cb02916877e74416e462934dae0a44685d498ad87672
```

```
$ docker inspect --format='{{.Id}}' arcsightsecurity/mgmt:1.10.0
```

```
sha256:fb046bb76298172eafc9ad54c6c3b7962d94539f2fef89ee140ef7428fe7f242
```

```
$ docker inspect --format='{{.Id}}' arcsightsecurity/rethinkdb:1.10.0
```

```
sha256:fca3b424936f8f3e3e0297b6d90ae38e2dac9aa7042ae6460163c4e574c50c94
```

```
$ docker inspect --format='{{.Id}}' arcsightsecurity/search-engine:
1.10.0
```

```
sha256:5a10f0ceba8d77fa0678d98b582c19e23efc638d5af80d28481775f3a213476f
```

What's next

"Generating a key pair on the Vertica cluster node 1 and installing the Investigate Vertica database" below

Generating a key pair on the Vertica cluster node 1 and installing the Investigate Vertica database

About

In a Vertica cluster deployment, generate a key pair on node 1 and then copy the public key to all nodes, including node 1. This enables password-less SSH access from the node 1 server to all the other node servers in the cluster.

The following is an example of enabling password-less SSH. For additional examples, see http://www.linuxproblem.org/art_9.html

Note:

- LVM partitioning is not supported. See Operating System Configuration Task Overview.
- ext3 or ext4 formatting is supported. xfs is not supported. See [Supported file systems in the Vertica 8.0.x documentation](#).
- Review the licensing requirements described in [Managing Licenses](#).

Procedure

Prerequisite

- Review the system requirements for ArcSight Investigate Vertica Database and proxy settings as outlined in ArcSight Investigate and Event Broker prerequisites.
- Disable the firewall.

To enable the firewall use the following:

<https://my.vertica.com/docs/8.0.x/HTML/index.htm#Authoring/InstallationGuide/BeforeYouInstall/EnsurePortsAreAvailable.htm>

1. On the node 1 server, run the `ssh-keygen` command on the node 1 server.

Example:

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

2. Copy the key from node 1 to all nodes, including node 1, using the node's IP address.

Example:

```
ssh-copy-id -i ~/.ssh/id_rsa.pub root@11.111.111.111
```

The system displays the key fingerprint and requests to authenticate with the node server.

3. Enter the required credentials of the nodes.

The operation is successful when the system displays the following message:

```
Number of key(s) added: 1
```

4. To verify that the key was successfully installed on the node, run the following command from node 1 to the target node to verify that it can successfully log into the node.

```
ssh 'root@11.111.111.111'
```

5. Repeat steps 1 through 4 for all nodes.

6. On the node 1 server, create a folder for ArcSight Investigate Vertica Database:

```
mkdir /root/install-vertica/
```

7. Copy the ArcSight Investigate Vertica Database scripts:


```
arcsight-investigate-vertica-scripts.136ef34.tar.gz to /root/install-vertica
```

```
arcsight-investigate-vertica-scripts.136ef34.tar.gz to /root/install-vertica/
```

8. Verify that the tarball matches the md5 checksum:

```
cd /root/install-vertica
```

```
md5sum arcsight-investigate-vertica-scripts.136ef34.tar.gz
```

```
cat arcsight-investigate-vertica-scripts.136ef34.tar.gz.md5
```

Both outputs should match.

9. Extract the tar file:

```
tar xvfz arcsight-investigate-vertica-scripts.136ef34.tar.gz
```

10. Open the `vertica.properties` file `cat vertica.properties` and edit the following values:

Property	Value
<code>ssh_private_key</code>	<code>/root/.ssh/id_rsa</code>
<code>timezone</code>	Your timezone according to the format in: <code>/usr/share/zoneinfo/</code> of Linux systems. For example, <code>US/Pacific</code> , <code>Europe/Prague</code> , <code>Japan</code> . "UTC" is the default setting.
<code>hosts</code>	A comma separated list of the ArcSight Investigate Vertica Database servers in IPv4 format (1.1.1.1)
<code>license</code>	Download the license file from the Software Entitlements portal. Place the license file on your filesystem, and then point to this file from license parameter.
<code>rpm</code>	<code>/root/install-vertica/data/vertica-8.0.1-5.x86_64.RHEL6.rpm</code>
<code>dba_user</code>	<dbadmin> Encrypted value provided during Vertica installation
<code>database</code>	<code>investigate</code>
<code>dbpassword</code>	<dbadmin> Encrypted value provided during Vertica installation
<code>ssl_enable=1</code>	Use this option if your database supports an SSL connection
<code>use_p2p=1</code>	Use this option in case your infrastructure does not support broadcast messaging or your nodes are not located on the same subnet. You should also use this option for all virtual environment installations, regardless of whether the virtual servers are on the same subnet or not.

11. Run the installation and create schema script.

The first command installs Vertica. The second one creates the schema "investigation" and adds a table "events".

```
./vertica_installer install
```

```
[root@n15-214-137-h78 install-vertica]# ./vertica_installer install
```

Please specify a username for the DB Admin user :

```
dbadmin
# =====
Username added successfully
# =====
Please specify a password for the DB Admin user :
Re-enter password:
# =====
Password added successfully
:
:
./vertica_installer create-schema
After installation completes, safeguard your database admin credentials.
```

12. Run the ArcSight Investigate Vertica Database scheduler (Kafka):

```
./kafka_scheduler create <EB Master Node IP>:39092 number_of_partitions
```

Note: number_of_partitions default is 6. However, it should match with the number defined in Event Broker master file, /opt/arcsight/installer/installer.properties:

```
predeploy.eb.init.noOfTopicPartitions
```

13. The Kafka Scheduler supports the following commands:

Action	Command	Description
Stop	./kafka_scheduler stop	Stops all running scheduler instances
Start	./kafka_scheduler start	Starts scheduler for all Kafka instances registered after performing a stop operation first.
Add host	./kafka_scheduler add host1:9092,host2:9092	Adds new hosts to running Vertica cluster. Add one by one, not comma-separated list, for additional hosts you want to add to the Vertica cluster. ./kafka_scheduler create <EB Master Node IP>:9092
Create	kafka_scheduler create <EB Master Node IP>:39092 # Default number is applied ./kafka_scheduler create <EB Master Node IP>:39092 <number_of_partitions>	Creates a new Kafka scheduler
Status	kafka_scheduler status host1:9092	Presents the status of running Kafka scheduler including count of imported/rejected messages
Delete meta data	kafka_scheduler delete	Deletes the meta data. After doing this, immediately run the kafka_scheduler create command.

14. Run `./kafka_scheduler status` to ensure that setup finished.

If the error message, "ERROR 4568: Relation investigation.rejected_events" appears, check the file, `/root/install_vertica/kafka_scheduler.log` for details.

Based on the log information, see the Release Notes for a workaround.

See Also

[ArcSight Investigate deployment troubleshooting and FAQs](#)

What's next

[Install ArcSight Investigate and Event Broker using the ArcSight Installer](#)

Configure ArcSight Investigate components

Configuring the ArcSight Investigate Vertica database connection in the ArcSight Installer

Procedure

1. Click **Configuration > ArcSight Investigate > Vertica**.
2. Enter the following information and then click **Save**:
 - Vertica host — Vertica IP
 - Vertica user name — dbadmin
 - Vertica database — Investigate
 - Vertica password — dbadmin

Configuring TLS on the ArcSight Investigate Vertica database in the Vertica server

About

The components of the ArcSight Investigate system interact using encrypted communication implemented using the Transport Layer Security (TLS) cryptographic protocol. The components managed by the ArcSight Installer are deployed with encrypted communication. This procedure provides instructions for distributing the key and certificate files on all ArcSight Investigate Vertica Database nodes and enabling TLS on the database.

Requirements

- A valid digital certificate signed by a certificate authority (CA). This includes two files:
 - Server certificate file (server.crt)
 - Root certificate file (root.crt)
- Private key file (server.key)

Note: The database does not need to be running when you distribute the key and certificate files.

Procedure

1. Copy the .crt and .key files to one of the ArcSight Investigate Vertica Database nodes.
2. Run the Vertica Administration Tools, as described in Using the Administration Tools in the Vertica documentation.
3. From the Main Menu, select **Configuration Menu** and click **OK**.
4. In the **Configuration Menu** screen, select **Distribute Config Files** and click **OK**.
5. In the **Select a category of files to copy** screen, select **SSL Keys** and click **OK**.
6. In the **Select database** screen, select the database on which you want to distribute the files and click **OK**.
7. In the **Select files to install** screen, modify the file path to the location to which you copied the files and click **OK**.

The names of the files should be:

- server.crt
- server.key
- root.crt

8. Run the Administration Tools again.

In the Main Menu screen, select **Connect to Database** and click **OK**.

9. When prompted, enter the database password.
10. Run the following command: `ALTER DATABASE mydb SET EnableSSL = 1;`
11. Restart the database.

Configuring Vertica SSL

About

The ArcSight Installer contains the script, `/opt/arcsight/installer/k8s/master/cert-utils.sh` which provides a tool that enables you to generate a certificate signed by the root certificate authority used by Kubernetes and all modules.

Procedure

1. Connect to the master node (where installation were run) and run `./cert-utils.sh generate-certificate vertica => script produce vertica.key and vertica.crt`
You can change `vertica` to your host name.
2. Copy `vertica.key` and `vertica.crt` to all Vertica nodes.
It is also needed to copy there certificate of certificate authority (default `/opt/arcsight/kuberntes/ssl/ca.crt`)
3. On each node run the `su - -c adminTools dbadmin => vertica Administration Tool`.
Use the user specified in Vertica configuration.

- a. From the **Main Menu** in the **Administration Tools**, select **Configuration Menu**, and then click **OK**.
- b. From the **Configuration Menu**, select **Distribute Config Files** and then click **OK**.
- c. Select **SSL Keys** and then click **OK**.
- d. Select the database on which you want to distribute the files (the database from configuration), and then click **OK**.
- e. Add the file locations for the `vertica.crt`, `vertica.key` and `ca.crt` (certificate authority certificate) files, and then click **OK** to distribute the files.

See Also

<https://my.vertica.com/docs/8.0.x/HTML/index.htm#Authoring/Security/SSL/ConfiguringSSL.htm>

Configuring the SMTP server in ArcSight Installer

About

Configure access to your SMTP server in ArcSight Installer to enable users that you create in ArcSight Investigate to receive notification emails.

Procedure

Location: ArcSight Installer

1. Go to **Configuration > ArcSight Investigate** and then click the **User Management** tab.
2. In the **User Management** tab, enter the following information and then click **Save**:
 - SMTP Host
 - SMTP Port
 - SMTP User Name
 - SMTP Password
 - Sender Address

Configure session and search settings in ArcSight Installer

About

You can configure the following properties:

- Session timeout

When the user session ends, the user is redirected to the login screen in order to log in again. The default session timeout is 60 minutes.

- Search query timeout

Search queries may take a long time and impact performance. You can put a limitation on the amount of time a search query runs. The default search query timeout is 60 minutes.

Procedure

Location: ArcSight Installer

1. Click **Configuration > ArcSight Investigate**, and then click the **General** tab.
2. In the **General** tab, do the following and click **Save**:
 - In the **Session timeout** field, enter the maximum time (in minutes) that you want a session to run.
 - In the **Search query timeout** field, enter the maximum time (in minutes) that you want a search query to run.

Configure Event Broker for management by ArcMC

- Ensure that the Event Broker host can be identified using an IP address
- Register the Event Broker host with ArcMC
- Manage nodes in ArcMC

See "Adding a host" in *ArcMC 2.6 Beta Administrator's Guide*.

Configure SmartConnectors

- Set Event Broker destinations for your topology in the SmartConnector Configuration screen (see the *SmartConnector User's Guide*).
- Make sure events are coming in and returned in the Search function.
- Add more connectors as desired—and verify.

Establishing the system admin

About

When you log in to ArcSight Investigate for the first time, you need to create the first user in the system. This user is assigned the system admin role.

Procedure

1. Open `https://master-ip`
2. From the Welcome page, enter the name, email, and password information for the system admin and then click **Create System Admin**.
3. From the Login page, enter the credentials for the system admin.

What's next

[Adding Kubernetes worker nodes](#)

Generate signed certificates for consumers

Event Broker consumers need a signed certificate from the Event Broker to establish secure communication.

There are a number of methods for generating signed certificates. Each used for different use cases:

- ArcSight Installer includes a utility for generating a signed certificate from the CA that is configured in the system (either Kubernetes or another trusted CA). You can use this utility to generate certificates for other components in the system, such as the ArcSight Investigate Vertica database.
- ArcSight Installer includes a utility for generating a signed certificate from a Certificate Signing Request (CSR) file. You can use this utility to generate certificates for client authentication, such as ESM and Logger.

Note: You can perform these procedures only after Kubernetes is installed.

Generate a signed certificate from the system CA

Procedure

1. Run the following command on the Kubernetes master server:
`/opt/arcsight/installer/k8s/master/cert-utils.sh generate-certificate`
The following argument is required:
 - `$2`: FQDN - fully qualified domain server nameThe following files are created in the directory where you ran the command:
 - `<FQDN>.crt`
 - `<FQDN>.key`
2. Copy the files to the server for which you generated the certificate.

Generating a signed certificate from a CSR

Procedure

1. Copy the CSR file to the Kubernetes master server.
2. Run the following command on the Kubernetes master server:

```
/opt/arcsight/installer/k8s/master/cert-utils.sh sign-certificate-request
```

The following arguments are required:
 - \$1: The CSR file (full path).
 - \$2: The name of the CRT you want to create (without the crt extension).A certificate and a private key are created
3. Copy the files to the server for which you generated the certificate.

Uninstalling ArcSight Investigate

About

Uninstalling ArcSight Investigate requires two basic steps:

- Uninstall ArcSight Installer.
- Uninstall Kubernetes.

Procedure

1. Uninstall ArcSight Installer. Run the following commands on the master server in this order:

```
yum erase -y arcsight-installer-1.10.7-ga110.x86_64.rpm
```
2. Uninstall Kubernetes.
Run the following command on all the worker nodes and on the master server and then reboot:

```
/opt/arcsight/kubernetes/uninstall.sh
```

```
yes
```

```
yes
```

The system reboots automatically.
3. After the server reboots, run on the master node:

```
rm -rf /root/.kube /opt/arcsight
```
4. On the worker node(s), run to do cleanup:

```
rm -rf /root/.kube /opt/arcsight /root/HPESW_ITOM_Suite_Platform_Worker
```

Caution: Do not re-install Event Broker/Investigate without first deleting the data files. If not deleted, ArcSight/Investigate will use old data files and not function properly.

See Also

[ArcSight Investigate deployment troubleshooting and FAQs](#)

ArcSight Investigate deployment troubleshooting and FAQs

Troubleshooting

Where to find the logs

To troubleshoot issues, capture the following logs. Logs are found under the pod number.

- zookeeper_container.log
- kafka_container.log
- schema-registry_container.log
- webservice_container.log

Pod starting order

After deploying Event Broker, 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). 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, Kafka Manager
5. Transformation Stream Processor, Routing Stream Processor

Cannot query zookeeper

Symptom: when you run `kubectl get pods` command to get status of the pods and you see that downstream pods (see the pod start order) do not stay up and the status is a 'CrashLoop'-type error.

Conditions to look for:

- 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 `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: Cannot elect a leader. If you see this type of error, check the check conditions above.
- Socket errors: 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

- You see indicating that broker Cannot Register ID: It might cause by multiple brokers with the same ID. This is a rare situation that can occur when you are add and removing nodes from the cluster and you do not define the cluster properly.

How to verify whether this is an issue: Connect to each system that is running a Kafka broker and check the assigned `broker.id` value of each. The `broker.id` value defined on each Kafka node must be unique.

```
# cat /opt/arcSight/k8s-hostpath-volume/eb/kafka/meta.properties
version=0
broker.id=1001
```

SSL Connection Error

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

- Cannot communicate with other brokers: If you see this type of error, host names may not be configured properly. It is possible that the node cannot perform reverse look up or that DNS is not set up properly.
- Out of disk condition (logs or data): <need to provide the actual error>: It is important to configure retention size correctly so that this does not occur. If the error occurs on one node, and other nodes are operating with no problems, you can delete the files (the other nodes in the cluster will continue to handle the event flow), change retention size on the topic manually, then restart the pod.

One option is to add more nodes to the cluster. Each node will store less data. Contact support or services if you want to take this approach.

Example: A Kafka node goes down, then the Schema Registry goes down. The Schema Registry will not restart while one Kafka node is down since it requires that all Kafka nodes to be up.

- Work to get the Kafka node back up. the Schema Registry pod requires that all Kafka pods are running. The Schema Registry will continue to check whether its dependency pods are running until it they are met. If Schema Registry crashes while performing this check, then K8s will restart the pod and Schema Registry will continue to check.
- Consequence: The Transformation Stream Processor (C2AV) will not function while SR is down. If SR is down for an extended period of time, the message queue in the EB-CEF topic increase as because messages are not being processed (converted to Avro). As long as the topic retention policy is set large enough, this should cause a problem. The message queue for the CEF topic will continue to increase without events being deleted. If the topic size or time range reaches the retention policy, then older messages will be deleted from the topic.
- If the Kafka node is permanently down, you might be able to edit the YAML file for Schema Registry. (This is an advanced task. It must be done by services or support.)

The Event Broker pre-defined topics are not created on the initial Event Broker first deployment

- Symptom: the Bootstrap Web Service log contains 500 response code (the response from SR), and topics are not created.
- Work around: Undeploy Event Broker containers, and then redeploy them.

One or more connectors cannot send data to Kafka

- Check whether the connection configuration is set properly in the connector.
- Check that the encryption mode (TLS, TLS+FIPS, TLS+CA, TLS+FIPS+CA) is the same for the Connector and Event Broker.
- Make sure you can connect to the Kafka port on the system and that there are no network issues.
- If you encounter a certificate error ("Cannot retrieve the certificate") when connecting, ensure 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. You need to fix the broker issue to that it comes up. In general, topics should be configured with replication > 1 so to prevent this scenario.
- Kafka is resyncing: This may cause event throughput slowdown, but will not stop event flow.
- Check whether the disk full.

A consumer cannot read events from Kafka

- New set up: Vertica Kafka scheduler: Check that kafka scheduler is configured to communicate to Kafka port 39092.
- Working at first, but stopped working: Offset is not 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 meta data. After doing this, immediately run the Kafka_scheduler create command to set up the scheduler.
- New set up: Check the network connection.
- New set up and existing set up: Check whether the broker is down.
- Existing set up: you did not configure all brokers that contain the topic the consumer connects to, and the brokers which are configured in that consumer are down.
- New set up: If you are encountering SSL connection related errors, check the steps that you used to import certificates to both Event Broker and consumers.

An Event Broker component crashes: ArcMC Rest API, stream processors (Routing and Transform)

- At Start Up: Check the container start up order (above). Have any of the dependency pods not started or crashed?
- Memory: JVMs require more memory than the system has available.
- All: check whether there are too many open sockets.

Event Broker EPS is lower than expected

- Check whether there are resource constraints on brokers: CPU, memory, disk is full. Check usage at system level (or with ArcMC).
- Network bottleneck.
- Stream processor is not able to keep up with transformation; SP is constrained in some way (resources). In ArcMC, the Stream Processor metric will be lower than the connector EPS.
Check that you have sufficient resources
- Expected file system
- Memory
- CPU
- File descriptor loads

FAQs

Which pods in Kubernetes comprise the Event Broker deployment?

- Event Broker pods
 - confluentinc pods: kafka; schemaregistry; zookeeper
 - Event Broker pods: c2av-processor, kafka-manager; orches; routing-processor; web-service

Related topic: [ArcSight Investigate and Event Broker prerequisites](#)

Which pods in Kubernetes comprise the ArcSight Investigate deployment?

- Hercules pods: management, proxy, rethinkdb, search

Related topic: [ArcSight Investigate and Event Broker prerequisites](#)

How can I find out more about Kafka and Apache Zookeeper?

See these resources for more about Kafka and Apache Zookeeper.

- Kafka: <https://sookocheff.com/post/kafka/kafka-in-a-nutshell>
- [Benchmarking Apache Kafka: 2 Million Writes Per Second \(On Three Cheap Machines\) | LinkedIn Engineering](#)
- [How to choose the number of topics/partitions in a Kafka cluster? - Confluent](#)
- [Apache Kafka](#)
- [Introduction to Kafka and Zookeeper](#)
- [Introduction to Apache ZooKeeper | Apache ZooKeeper Tutorials Setting up Apache ZooKeeper Cluster | Apache ZooKeeper Tutorials](#)

Can I use my existing Event Broker v1.0 with ArcSight Investigate + Vertica?

No. ArcSight Investigate requires Event Broker 2.0. You can migrate your data from Event Broker 1.0 using the Event Broker Data Migration utility. Check with ArcSight Support about the availability of this tool.

Related topic: *Event Broker Data Migration Tech Note*.

With only a single install method for ArcSight Investigate and Event Broker, how can I distribute each of these (pods) across Kubernetes worker nodes?

The ArcSight Installer application automatically manages the distribution and deployment of Kubernetes worker nodes.

Related topic: [Adding Kubernetes worker nodes](#)

How do I check the hostname of machines on which Event Broker is installed?

Example:

```
[root@n11.222.444.h11 ~]# kubectl get node -L fdqn
NAME                STATUS    AGE           FDQN
11.222.333.222      Ready    1d            n11.222.333.h222.domainname.com
11.222.444.11       Ready    1d            n11.222.444.h11.domainname.com
```

After checking labels on the above Event Broker, it shows both master and work node are using the same label. Need to make sure they are correct Label.

To verify whether the hostname is set correctly:

```
# ssh root@11.222.777.111 "hostname -f"
root@11.222.777.111's password:
n11.222.777.h111.domainname.com
# ssh root@11.222.777.111 "nslookup 11.222.777.111 | grep 'name =' "
root@11.222.777.111's password:
111.333.222.11.in-addr.arpa      name = n11.222.777.n111.domainname.com.
```

Appendix A: Adding Kubernetes worker nodes

About

- The default deployment consists of a Kubernetes master node, and three Kubernetes (K8s) worker nodes: three worker nodes for ArcSight Event Broker, and one Kubernetes master node for ArcSight Investigate.
- You can add additional worker nodes to extend the Kafka cluster nodes.
- Once you add the new worker nodes, labels can be used to assign specific pods to them, like with Kafka.

Procedure

1. To add a new Kafka node add a new worker node and label it, update the petset replica count and update the installer properties.

For example, if you want to add two more nodes to an existing 3-node Kafka cluster to create a 5-node Kafka cluster:

```
/opt/arcsight/installer/k8s/node/install.sh -w <new_worker_node_1_IPv4_address>
```

```
/opt/arcsight/installer/k8s/node/install.sh -w <new_worker_node_2_IPv4_address>
```

```
kubectl label --overwrite node <new_worker_node_1_IPv4_address> kafka=yes
```

```
kubectl label --overwrite node <new_worker_node_2_IPv4_address> kafka=yes
```

```
kubectl scale petset eb-kafka --replicas=5
```

2. Update `installer.properties` with the new Kafka node count so future deploy/undeploy will have the correct Kafka node number.

```
$ vi /opt/arcsight/installer/installer.properties
```

```
...
```

```
predeploy.eb.kafka.count=5
```

```
...
```

```
$
```

3. Synchronize the data across all Kubernetes worker nodes.

From the Kafka manager (Cluster | event-broker | Topics), do the following:

- a. Click **Generate Partition Assignments**.
- b. Click **Confirm Assignments: Generate Partition Assignments**.
- c. Go to topic list.
- d. Operation: Select **Run Partition Assignments**.
- e. Run Assignments: Run Partition Assignments.
- f. Go to Reassign Partitions

- g. Wait until the Status|Completed shows a date/time of completion
- h. Go to Preferred Replica Election
 - i. Preferred Replica Election: Run Preferred Replica Election
- j. Wait until the Status|Completed shows a date/time of completion

See Also

- [ArcSight Investigate and Event Broker prerequisites](#) for more about `installer.properties`
- [ArcSight Investigate deployment troubleshooting and FAQs](#)

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 Deployment Guide (ArcSight Investigate 1.10)

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!