



Hewlett Packard
Enterprise

HPE Security ArcSight Investigate

Software Version: 2.01

Deployment Guide

December 8, 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.softwagrp.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 Technical Support Page: https://softwaresupport.softwagrp.com/support-contact-information
Support Web Site	https://softwaresupport.softwagrp.com/
ArcSight Product Documentation	https://community.softwagrp.com/t5/ArcSight-Product-Documentation/ctp/productdocs

Revision History

Date	Description
December 8, 2017	Initial release of this document.

Contents

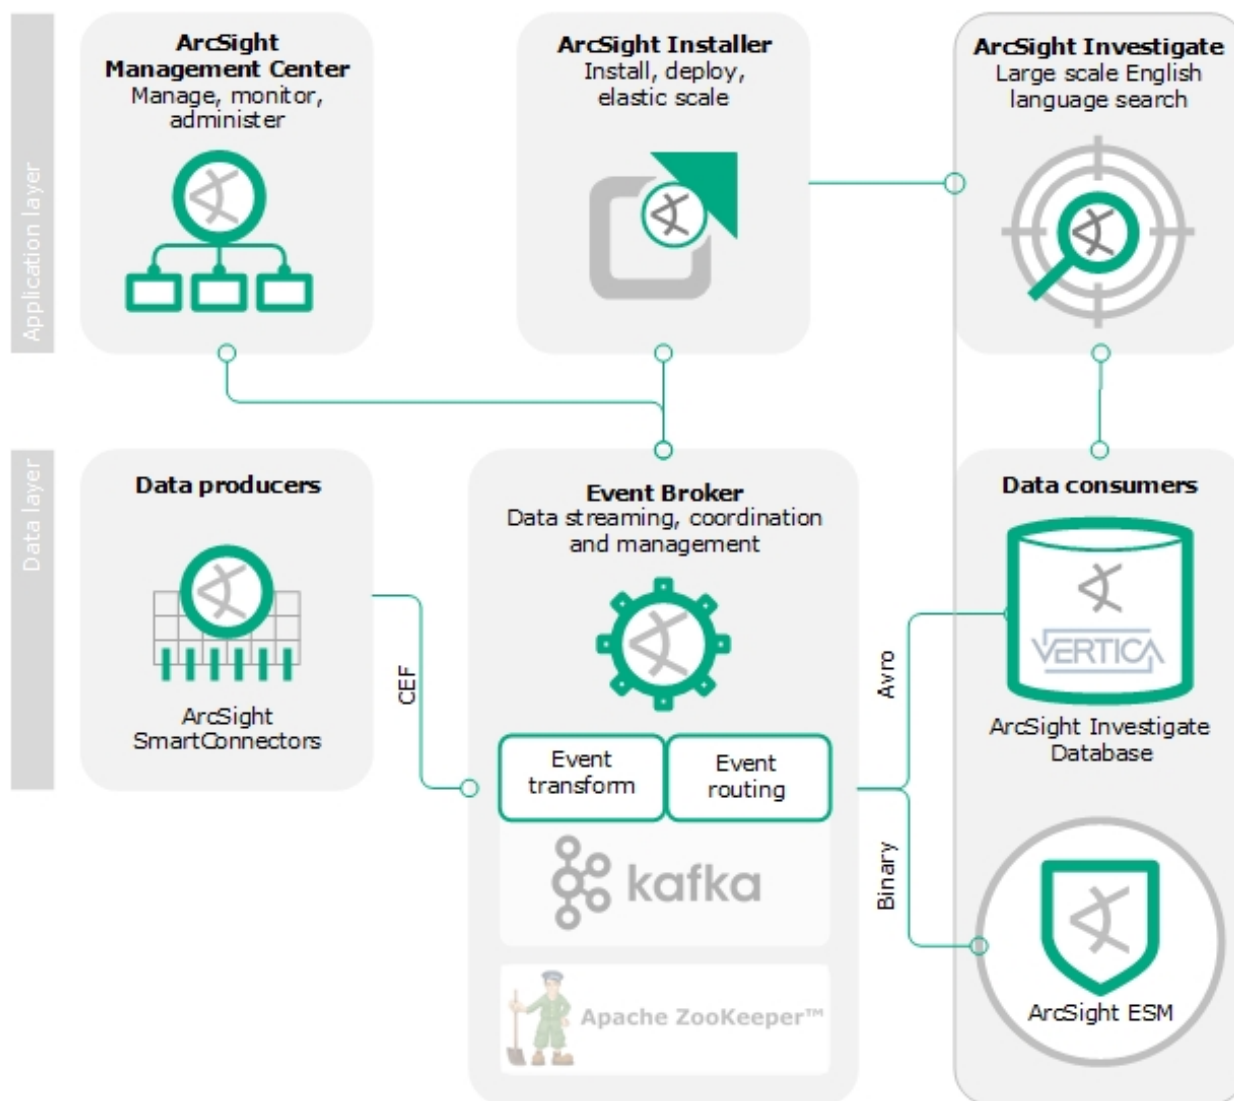
Chapter 1: Introduction	1
About ArcSight Investigate	1
ArcSight Investigate deployment architecture	2
Deployment overview	4
Planning your deployment	5
TLS planning	5
Network planning	6
Plan encryption modes	6
Chapter 2: ArcSight Investigate support matrix	8
Supported operating systems	8
Supported browsers	8
Supported product compatibility	8
Chapter 3: Prerequisites for installation	9
System requirements	9
General sizing guidelines	9
RPM requirements	11
Network requirements	11
Firewall requirements	11
Increasing per-user process limits	12
NTP requirements	12
Generating a key pair on the master node for worker nodes	13
Configuring proxy settings	13
Chapter 4: Install ArcSight Investigate	15
Labeling nodes	15
Obtaining ArcSight Investigate images (online)	15
Obtaining ArcSight Investigate images (offline)	16
ArcSight Installer tasks	16
Deploying ArcSight Investigate images	17

Configuring Event Broker for ArcSight Investigate	18
Chapter 5: Install Vertica	20
Generating the SSH key pair	20
Setting security-enhanced Linux (SELinux) to permissive	21
Disabling the firewall	21
Installing the ArcSight Investigate Vertica database	21
Chapter 6: Configure ArcSight Investigate and components	25
Establishing the system admin	25
Configuring the ArcSight Investigate Vertica database connection	25
Configuring the SMTP server	26
Configuring session and search settings in ArcSight Installer	26
Configuring TLS on the ArcSight Investigate Vertica database in the Vertica server	27
Configuring Vertica SSL	28
Chapter 7: Uninstalling ArcSight Investigate	30
Appendix A: ArcSight Investigate deployment troubleshooting and FAQs	31
Troubleshooting	31
Installing the ArcSight Installer Platform fails	31
Where to find the logs	31
Pod starting order	31
SSL connection error	31
kubectl command is returning refused or time-out connection	32
Vertica Scheduler unable to read events from Kafka	32
FAQs	32
Which pods in Kubernetes comprise the ArcSight Investigate deployment?	32
Can I use my existing Event Broker v1.0 with ArcSight Investigate + Vertica?	32
Send Documentation Feedback	34

Chapter 1: Introduction

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. Investigate indexes the events from your data source so you can view and search on them. You can use the English-like search language to generate results from which to create reports and visualizations.

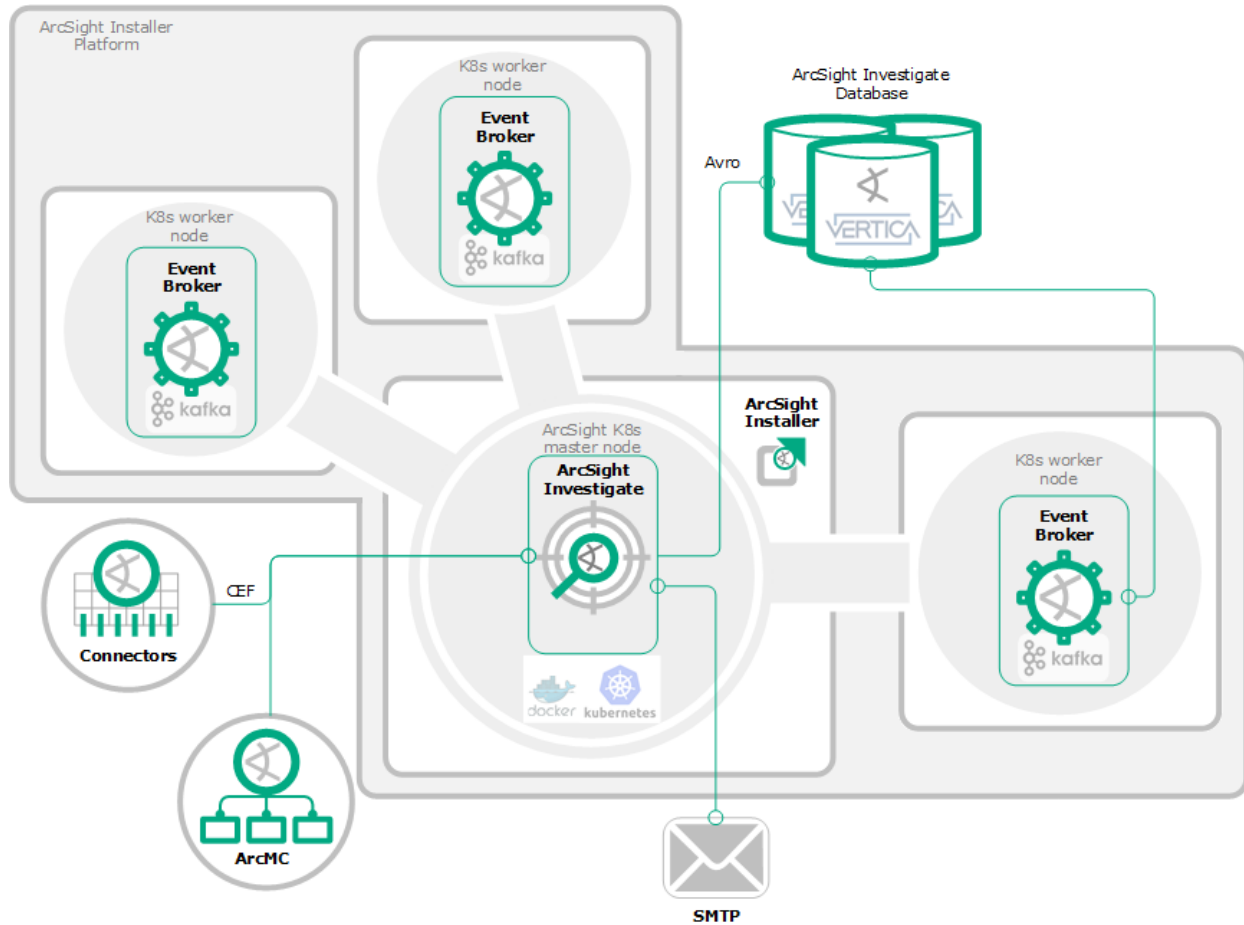


Component	Description
ArcSightInvestigate	High-capacity data management, search, and analysis web application.
ArcSight Installer	<p>A web application for deploying and configuring the ArcSight Investigate components, including Investigate and Event Broker.</p> <p>The components are managed in a Kubernetes cluster. The master node hosts the ArcSight Installer web application and the Investigate web application, and the worker nodes host the Event Broker.</p>
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.
ArcSightSmartConnectors	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. ArcSightSmartConnectors, installed and maintained separately, are producers that publish data to Event Broker. You can subscribe to data managed by Event Broker with Investigate, ArcSight Deployment Platform (ADP) Logger, ArcSight ESM, Apache HDFS, or your own third-party consumer.
Event Broker	ArcSightEvent Broker, a product of the ADP suite, 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 ADP.

ArcSight Investigate deployment architecture

ArcSight Investigate runs in Docker containers managed by Kubernetes and deployed from the ArcSight Installer application. The deployment typically consists of a Kubernetes master node, and three Kubernetes worker nodes: three nodes for ArcSight Event Broker and one node for ArcSight Investigate.

The image below is a typical representation of the deployment architecture.



Deployment component	Host	Functional contents
ArcSight Installer Platform	Install the platform on the master node and each work node.	ArcSight Installer application
Kubernetes master node	1 VM or physical server	<ul style="list-style-type: none"> Kubernetes master node Investigate ArcSight Installer application
Kubernetes worker nodes	3 VMs or physical servers	3 Event Broker nodes
Vertica database	3 physical servers	3 ArcSight Investigate Vertica database instances

Deployment component	Host	Functional contents
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.

Deployment overview

Before you can deploy ArcSight Investigate deployment you must first deploy the ArcSight Installer, Event Broker, and the Vertica database.

Note: ArcSight recommends installing and running these components in a test environment before putting them into production.

These components require configuration after you install Investigate containers.

1. Complete the installation requirements.
 - a. Ensure that you upgraded to ArcSight Installer 1.30 and Event Broker 2.11.

For how to upgrade Event Broker to version 2.11, refer to the *ArcSight Event Broker 2.11 Release Notes*. For how to install Event Broker 2.11, refer to the *ArcSight Event Broker 2.11 Deployment Guide*.
 - b. Ensure that Event Broker and Investigate each have a dedicated servers.

If other applications are running on the same server as Event Broker and Investigate, there will be a significant performance penalty and potential problems.
 - c. Generate an SSH certificate on the master node in order to allow connections to the worker nodes.

See ["Generating a key pair on the master node for worker nodes" on page 13](#).
2. Obtain the Investigate image.
 - For online retrieval, see ["Obtaining ArcSight Investigate images \(online\)" on page 15](#).
 - For offline retrieval, see ["Obtaining ArcSight Investigate images \(offline\)" on page 16](#).
3. Deploy Event Broker if necessary.

See the *ADP Event Broker Deployment Guide*.
4. Deploy Investigate.

See ["Deploying ArcSight Investigate images" on page 17](#).

5. Ensure that the Event Broker and Investigate images have completed deployment.

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

See ["Deploying ArcSight Investigate images" on page 17](#).

6. Install the Vertica database.

See ["Installing the ArcSight Investigate Vertica database" on page 21](#).

7. Configure Event Broker if necessary.

See the *ADP Event Broker Deployment Guide*.

8. Configure Investigate, including the Investigate Vertica database, SMTP server.).

See ["Configure ArcSight Investigate and components" on page 25](#)

Planning your deployment

Before deploying, ensure that you have the latest version of this document, available for download from the [ArcSight Product Documentation Community on Protect 724](#).

TLS planning

The various components in the ArcSight Investigate system interact using encrypted communication implemented using Transport Layer Security (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.

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 DNS 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. Define these variables per user, and not system-wide.

Plan encryption modes

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>
ArcSight SmartConnectors	<p>ArcSight SmartConnectors 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>	9093	<ul style="list-style-type: none">• TLS• FIPS• ClientAuth	<i>SmartConnector User Guide</i> <i>ArcMC Administrator's Guide</i>

Product	Preparations needed	Open ports	Supported encryption modes	Guidance documentation
ArcSight ESM (optional)	<p>ArcSight ESM can be installed and running prior to installing ArcSight Investigate and Event Broker.</p> <p>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.</p>	9093	<ul style="list-style-type: none"> • TLS • FIPS • ClientAuth 	<p><i>ESM Installation Guide</i></p> <p><i>ESM Administrator's Guide</i></p>
ArcSight Logger (optional)	ArcSight Logger can be installed and running prior to installing Event Broker.	9093	<ul style="list-style-type: none"> • TLS • FIPS • ClientAuth 	<p><i>Logger Administrator's Guide</i></p>

Chapter 2: ArcSight Investigate support matrix

Supported operating systems

Version	Component	Operating system
2.01	ArcSight Investigate	RHEL 7.3 64-bit* RHEL 7.4 64-bit CentOS 7.3 64-bit* CentOS 7.4 64-bit * Linux kernel version 3.10.0-514.26.2.el7.x86_64 (or above)
	ArcSight Investigate Vertica 8.1.1-3 database	RHEL 7.3 and CentOS 7.3

Supported browsers

Browser	Version
Microsoft Edge	Version available at the time of release.
Microsoft Internet Explorer	Version available at the time of release.
Google Chrome	Version available at the time of release.
Mozilla Firefox	Version available at the time of release.

Supported product compatibility

Product	Version
ArcSight Event Broker	2.11
ArcSight SmartConnector	7.5 and later
ArcMC	2.7.1
ArcSight Logger	6.5.1
ArcSight ESM	6.11

Chapter 3: Prerequisites for installation

- Ensure that Event Broker and Investigate have a dedicated server.
If other applications are running on the same server as Event Broker and Investigate, there will be a performance penalty and potential problems.
- Ensure that your file system type is ext4.

System requirements

General sizing guidelines

Provision the servers (or VMs) that you are using for the deployment, based on the general sizing guidelines provided here. This information is based on a default setup.

For tailored sizing recommendations for a production environment, contact ArcSight Customer Support.

For supported platforms and operating systems, see the ArcSightInvestigate Support Matrix.

Component	Nodes	Resources needed	Needed ports
ArcSight Investigate + Event Broker	1 master 3 worker	<ul style="list-style-type: none"> One CPU with 24 cores 32 GB RAM 8 TB disk space Linux kernel version 3.10.0-514.26.2 (or above) 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>Note: If you choose to deploy ArcSight Investigate on a worker node, the nginx reverse proxy used to connect to Investigate is always deployed on the master node. Therefore, no matter where Investigate is deployed in a Kubernetes cluster, you should always access Investigate using the host/IP of the master node.</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> <p>Investigate: 5443, 21085, 30001</p>
Investigate Vertica database	3	<p>Important: The Vertica database must be installed on the same sub-network as the 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 8.1.x Documentation for more information.</p>	5433
Vertica scheduler			
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>	

RPM requirements

The following packages need to be installed using rpm/yum on all systems (master and workers):

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

Java 8 needs to be available and accessible on all servers.

Network requirements

Caution: If the default network ranges specified here are in use in your network environment, the installation may fail, or random failures may be experienced after installation.

By default, the Installer uses the following network ranges:

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

For the /16 and /24 address ranges, ensure that your network is conflict free. If these address ranges are occupied and/or not accessible due to network configuration, utilize another address range by making corresponding changes to the POD_CIDR, SERVICE_CIDR and DNS_SVC_IP parameters in the `./<path to the secure location on master node>/arcsight-installer-<version>/arcsight-installer-master.sh` script.

Firewall requirements

The following ports need to be free and available for firewall configuration.

- **Kubernetes:** 2379,2380,3000,4001,4194,5000,5443,8080,8088,8200,8285,8443,10248-10252,10255
- **NFS:** 111,2049,20048,37189
- **Investigate:** 5443,21085,30001

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

Increasing per-user process limits

Procedure

1. Do the following on every Vertica and Kubernetes node.
Open the file `/etc/security/limits.d/20-nproc.conf`.
If you do not already have a `/etc/security/limits.d/20-nproc.conf` file, create one (and the `limits.d` directory, if necessary).
2. Add the lines below, 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 Kubernetes master and worker nodes, and Vertica nodes.
Nodes can be rebooted in any order.
4. Verify that all nodes are up and running by running the following command.

```
ulimit -a
```

NTP requirements

About

chrony is a versatile implementation of the Network Time Protocol (NTP). **chrony** keeps the system clocks of each of the cluster nodes 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 your system, install it with the following procedure.

Procedure

1. `yum install chrony`
2. Start **chronyd** to start and enable the **chrony** daemon.

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


3. Verify that **chrony** is operating correctly.

`chronyc tracking`

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 sudo 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.

Configuring proxy settings

About

Comment out proxy data in the `/etc/profile.d/proxy.sh` file on all nodes, if it is being used.

If you are using a proxy server in your environment, then add your proxy data to the `~/.bashrc` file.

Procedure

Update the `.bashrc` file according to the following example:

```
export http_proxy=http://<proxyserver>:8080/
export https_proxy=http://<proxyserver>:8080/
export HTTP_PROXY=http://<proxyserver>:8080/
export HTTPS_PROXY=http://<proxyserver>:8080/

export no_proxy="<master ip>,<worker-1 ip>,<worker-2 ip>,<worker-3 ip>,localhost,<domain>"

export NO_PROXY="<master ip>,<worker-1 ip>,<worker-2 ip>,<worker-3 ip>,localhost,<domain>"
```

Chapter 4: Install ArcSight Investigate

The ArcSight Installer is used to install ArcSight Investigate and Event Broker. Before running the ArcSight Installer, verify that you have set up the receiving systems according to guidelines in Investigate and Event Broker prerequisites. Multi-master installation is not supported.

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

This user guidance 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.

Labeling nodes

About

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

Procedure

SSH to the master node and label all nodes.

- `kubectl label --overwrite node {masternode_ip} investigate=yes`
- `kubectl label --overwrite node {workernode1_ip} kafka=yes zk=yes`
- `kubectl label --overwrite node {workernode2_ip} kafka=yes zk=yes`
- `kubectl label --overwrite node {workernode3_ip} kafka=yes zk=yes`

Obtaining ArcSight Investigate images (online)

About

Download the Investigate images .

Procedure

1. Obtain Investigate.
 - a. Download Investigate images.

```
cd /opt/arcsight/kubernetes/scripts  
./downloadimages.sh --suite investigate -r docker -o arcsightsecurity
```
 - b. Pick the 2.01 version.
 - c. Upload the images to the local Docker registry.

```
./uploadimages.sh --suite investigate
```
2. Obtain Event Broker if not already installed.

See the *ADP Event Broker Deployment Guide*.

Obtaining ArcSight Investigate images (offline)

About

Obtain the Investigate images tar files from the ArcSight FTP server.

Procedure

Investigate

1. Download the `arcsight-investigate-*.tar` file.
2. Place the `arcsight-investigate-*.tar` file in `master:<offline install directory>`.
3. Upload Investigate images.

```
cd <offline install directory>  
tar xvf arcsight-investigate-*.tar
```

All Investigate related images will be stored in the `./investigate` directory.

```
cd/opt/arcsight/kubernetes/scripts  
./uploadimages.sh -s investigate -d <offline install  
directory>/investigate
```

Event Broker

If not already deployed, obtain Event Broker images.

See the *ADP Event Broker Deployment Guide*.

ArcSight Installer tasks

From the ArcSight Installer (UI page), you can check the status of the master and worker nodes, deploy the Investigate images, and configure Investigate .

Deploying ArcSight Investigate images

About

With your browser pointed at the ArcSight Installer, the Investigate deployment option is located in the Deployment page, ready for deployment, with an initial status of **OFF**.

Procedure

Location: ArcSight Installer > left navigation > Deployment > Deployment page

1. Login to the ArcSight Installer (UI page).
`https://<master-FQDN>:5443`
2. Click **Deploy** for ArcSight Investigate, then click 2.01.
A green circle containing a check mark appears. This indicates that the deployment is in progress.
3. To check the deployment status, run the command, `kubect1 get pods --all-namespaces` on the master node.

Although the status indicator in ArcSight Installer shows that the deployment is complete, the Investigate containers may be not be in a running state.

When deployment is complete, all pods under the `arcsightinvestigate1` namespaceshould be in the running state, as shown in the example below.

Note: It may take 2-5 minutes for all pods to start running.

NAMESPACE	NAME	READY	STATUS	RESTARTS	AGE
arcsighteventbroker1	eb-c2av-processor-0	1/1	Running	0	1d
arcsighteventbroker1	eb-kafka-0	1/1	Running	1	1d
arcsighteventbroker1	eb-kafka-1	1/1	Running	0	1d
arcsighteventbroker1	eb-kafka-2	1/1	Running	0	1d
arcsighteventbroker1	eb-kafka-manager-3844815475-p3fnd	1/1	Running	0	1d
arcsighteventbroker1	eb-routing-processor-0	1/1	Running	0	1d
arcsighteventbroker1	eb-schemaregistry-51771507-gv7mv	1/1	Running	1	1d
arcsighteventbroker1	eb-web-service-1189059977-c08vc	2/2	Running	0	1d
arcsighteventbroker1	eb-zookeeper-0	1/1	Running	0	1d
arcsighteventbroker1	eb-zookeeper-1	1/1	Running	0	1d
arcsighteventbroker1	eb-zookeeper-2	1/1	Running	0	1d
arcsighteventbroker1	suite-reconf-pod-eventbroker-gpqq6	2/2	Running	0	1d
arcsightinvestigate1	hercules-management-688604836-40jcl	2/2	Running	0	1d
arcsightinvestigate1	hercules-rethinkdb-0	1/1	Running	0	1d
arcsightinvestigate1	hercules-search-3729025617-kjndw	3/3	Running	0	1d
arcsightinvestigate1	nginx-ingress-controller-3790412081-bw5t3	1/1	Running	0	1d
arcsightinvestigate1	suite-reconf-pod-investigate-7w306	2/2	Running	0	1d
core	apiserver-15.214.134.8	1/1	Running	0	2d
core	arcsight-controller-1434481503-gja73	2/2	Running	0	2d
core	controller-15.214.134.8	1/1	Running	0	2d
core	heapster-apiserver-4160107731-qf1dc	1/1	Running	0	2d
core	ide-4211554324-dfzpw	2/2	Running	0	2d
core	ide-4211554324-naggo	2/2	Running	0	2d
core	ide-guestgw-1-1339164482-5jba9	2/2	Running	0	2d
core	kube-dns-3414883154-0hghg	3/3	Running	6	2d
core	kube-proxy-15.214.134.144	1/1	Running	0	2d
core	kube-proxy-15.214.134.144	1/1	Running	0	2d
core	kube-proxy-15.214.134.8	1/1	Running	0	2d
core	kube-proxy-15.214.137.14	1/1	Running	0	2d
core	kube-registry-810561521-93fem	1/1	Running	0	2d
core	kube-registry-proxy-412fg	1/1	Running	0	2d
core	kube-registry-proxy-41gbe	1/1	Running	0	2d
core	kube-registry-proxy-kalhd	1/1	Running	0	2d
core	kube-registry-proxy-qg014	1/1	Running	0	2d
core	kubesates-vaalt-94103177-9a3t3	1/1	Running	0	2d
core	mg-guest-1-1303047706-qphom	2/2	Running	1	2d
core	scheduler-15.214.134.8	1/1	Running	0	2d
core	suite-podf-pod-eventbroker	2/2	Running	0	1d
core	suite-db-314010429-w13q3	2/2	Running	0	2d
core	suite-controller-2922893421-3v1td	2/2	Running	0	2d
default	nginx-ingress-controller-zh7vx	1/1	Running	0	2d

- To undeploy ArcSight Investigate, click **Undeploy**.

Each time you deploy, make sure to reconfigure the Investigate Vertica database connection. The information does not persist when the application is undeployed.

Configuring Event Broker for ArcSight Investigate

About

Once you deploy Event Broker, you can then configure the Event Broker data pipeline for ArcSight Investigate from the ArcSight Installer.

Notes:

- Event Broker consumers need a signed certificate from the Event Broker to establish secure communication with Investigate (see the *ADP Event Broker Deployment Guide*).
- In the event of a planned redeployment of Event Broker without a restart of the cluster node systems, be sure to do a clean undeploy of event broker (see the *ADP Event Broker Deployment Guide*).

Procedure

Location: ArcSight Installer > left navigation > Configuration

1. Select **ArcSight Event Broker**.
2. Select **Replicas**.
3. Click **+** next to Transforming String Processor and click **Save**.

The number will change from 0 to 1.

Chapter 5: Install Vertica

Requirements for provisioning the Vertica server

- No LVM partition
- Partition type: ext4
- Minimum 2 GB swap space
- RHEL 7.3 or CentOS 7.3 only

Prerequisites for Installing Vertica

- Review the system requirements for ArcSight Investigate Vertica Database.
- Increase default user process limit by following the instructions in ["Increasing per-user process limits" on page 12](#)

Generating the SSH key pair

About

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.

Procedure

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

Example:

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

2. Copy the key from node 1 to all nodes, including node 1, using the node 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.

Setting security-enhanced Linux (SELinux) to permissive

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.

Disabling the firewall

About

This procedure is needed for the Vertica host only.

Procedure

1. 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
```

Installing the ArcSight Investigate Vertica database

Procedure

Prerequisite

Generate a key pair on the Vertica cluster node 1 (see ["Install Vertica" on page 20](#)).

1. On the node 1 server, create a folder for the ArcSight Investigate Vertica Database:
`mkdir /root/install-vertica/`
2. Copy the Investigate Vertica Database scripts:
`arcsight-investigate-vertica-scripts.<hash>.tar.gz` to `/root/install-vertica`
`arcsight-investigate-vertica-scripts.<hash>.tar.gz.md5`
3. Verify that the tarball matches the MD5 checksum:
`cd /root/install-vertica`
`md5sum arcsight-investigate-vertica-scripts.<hash>.tar.gz`
`cat arcsight-investigate-vertica-scripts.<hash>.tar.gz.md5`
Both outputs should match.
4. Extract the tar file:
`tar xvfz arcsight-investigate-vertica-scripts.<hash>.tar.gz`
5. Edit the `vertica.properties` file.
 - The `hosts` and `license` properties must be updated while the other properties are optional.

Property	Value
# environment settings	
ssh_private_key	/root/.ssh/id_rsa
timezone	Your timezone according to the format in: /usr/share/zoneinfo/ of Linux systems. For example, US/Pacific, Europe/Prague, Japan. "UTC" is the default setting.
# vertica settings	
# please be sure not to use loop-back addresses in case cluster would need to be performed.	
hosts	A comma separated list of the Investigate Vertica Database servers in IPv4 format (1.1.1.1)
license	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.
rpm	/root/install-vertica/data/vertica-8.0.1-5.x86_64.RHEL6.rpm
dba_user	<dbadmin> Encrypted value provided during Vertica installation
database	investigate
dbpassword	<dbadmin> Encrypted value provided during Vertica installation

Property	Value
ssl_enable=1	Use this option if your database supports an SSL connection
# Database users # DBAdmin	
dba_user	Database administrator
dba_password	Database administrator password
# search	
search_user	Search user for ArcSight Investigate
search_password	Search user password for ArcSight Investigate
## Tune DB Tuple Mover (TM) to for best ingestion performance, recommended: 4 active partitions, 5 threads for TM, 6000 MB for TM 5 threads for TM, 6000 MB for TM ## Use scripts/tuning_util.sh to modify below values after installing vertica	
active_partitions=4 tm_concurrency=5 #value in MB tm_memory=6000	Database tuning parameters
use_p2p=1	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.

- To install Vertica, run the following command:

```
./vertica_installer install
```

You will be prompted to set up two users, a database administrator and an Investigate search user. After installation completes, safeguard your database admin credentials.

- To create the schema and database tables, run the following command :

```
./vertica_installer create-schema
```

- To create a scheduler and related schema and tables, run the following command :

```
./kafka_scheduler create <EB Worker Node 1 IP>:9092,<EB Worker Node 2 IP>:9092,<EB Worker Node 3 IP>:9092 number_of_partitions
```

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

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

- The Kafka Scheduler supports the following commands:

Action	Command	Description
Stop	<code>./kafka_scheduler stop</code>	Stops all running scheduler instances
Start	<code>./kafka_scheduler start</code>	Starts scheduler for all Kafka instances registered after performing a stop operation first.
Create	<code>./kafka_scheduler create <EB Worker Node 1 IP>:9092,<EB Worker Node 2 IP>:9092,<EB Worker Node 3 IP>:9092 number_of_partitions</code>	Creates a new Kafka scheduler
Status	<code>kafka_scheduler status host1:9092</code>	Presents the status of running Kafka scheduler including count of imported/rejected messages
Delete	<code>kafka_scheduler delete</code>	Deletes the meta data. After doing this, immediately run the <code>kafka_scheduler create</code> command.

10. To check the Vertica status, run `./kafka_scheduler status`.

See Also

[ArcSight Investigate deployment troubleshooting and FAQs](#)

Chapter 6: Configure ArcSight Investigate and components

Once you deploy ArcSight Investigate, you can then configure the product from the Configuration page of the Installer. After changing a product setting, Investigate restarts. Wait until restart completes before logging into Investigate.

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.

Configuring the ArcSight Investigate Vertica database connection

Procedure

Location: ArcSight Installer

1. Click **Configuration > Investigate > Vertica**.
2. Enter the following information and then click **Save**:
 - Vertica host — Vertica node 1 IP
 - Vertica user name — See step 6 in ["Installing the ArcSight Investigate Vertica database" on page 21](#).
This is the Investigate search user name that you created during installation.
 - Vertica database — Investigate. This was defined during schema creation and cannot be changed.

- Vertica password — See step 6 in ["Installing the ArcSight Investigate Vertica database" on page 21](#)
This is the Investigate search user password.

Configuring the SMTP server

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

Configuring 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**.
2. From the **General** tab, do the following:
 - 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.
3. Click **Save**.

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/kubernetes/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>

Chapter 7: Uninstalling ArcSight Investigate

About

Uninstalling ArcSight Investigate requires two basic steps:

- Uninstall Kubernetes.
- Uninstall the Vertica database.

Procedure

1. Uninstall Kubernetes.
 - a. Run the following command on all the worker nodes and on the master node and then reboot:

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


yes
yes
 - b. After the server reboots, run the following command on the master node.

```
rm -rf /root/.kube
```


if /root/.docker exists,

```
rm -rf /root/.docker
```
 - c. On the worker node(s), run:

```
rm -rf /root/.kube /root/arcsight-installer-worker
```

Note: If you want to delete all your data as well, run the following command on the master node.

```
rm -rf /opt/arcsight /opt/kubernetes
```

Run the following command on the worker nodes:

```
rm -rf /opt/arcsight
```

2. Uninstall the Vertica database.

Run the following on all Vertica nodes:

- a. Stop the process.

```
kill -9 -f vertica #stop vertica process
```
- b. Remove the package.

```
rpm -e vertica-8.1.1-3.x86_64 #delete vertica package
```
- c. If you want to delete the configuration file used with your installation, you can choose to delete the /opt/vertica/ directory and all sub-directories using this command:

```
rm -rf /opt/vertica #delete Vertica data
```


See the [Vertica Analytics Platform Version 8.1.x Documentation](#).

Appendix A: ArcSight Investigate deployment troubleshooting and FAQs

Troubleshooting

Installing the ArcSight Installer Platform fails

Contact Technical Support.

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

SSL connection error

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

kubectl command is returning refused or time-out connection

If the `kubectl` command is returning refused or time-out connection, make sure the proxy is unset before repeating the command.

Vertica Scheduler unable to 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.

FAQs

Which pods in Kubernetes comprise the ArcSight Investigate deployment?

- Hercules pods: management, proxy, rethinkdb, search

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

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*.

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 (Investigate 2.01)

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!