



ArcSight Fusion 1.0 Administrator Guide

April 2020

Legal Notice

© Copyright 2020 Micro Focus or one of its affiliates.

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

For additional information, such as certification-related notices and trademarks, see <https://www.microfocus.com/about/legal/>.

Contents

About This Book	7
1 Understanding ArcSight Fusion	9
Understanding the Fusion Architecture	10
Understanding the Fusion Components	10
Part I Planning Installation and Deployment	11
2 Installation Checklist	13
3 Understanding the Installation Process	15
Components Needed for Deployment	15
ArcSight Fusion	15
Analytics Framework	15
Documentation	15
Methods of Deployment	16
Single-Node Deployment	16
Multi-Node Deployment	16
High Availability Deployment	16
Securing Communication Between Components	16
4 Installation Options	19
Installation Using Scripts	19
Manual Installation	19
Deciding to Use the Scripts or Manual Installation Method	19
5 Installing Fusion and ESM on the Same Node	21
Part II Installing CDF and Deploying Fusion	23
6 Preparing Your Environment	25
7 Installing ArcSight Fusion by Using Scripts	27
Prerequisites	27
Understanding the Installation Scripts	27
Using the Installation Scripts	27
8 Installing ArcSight Fusion Manually	29
Preparing Your Environment for CDF	29
Configuring the Nodes	29

Setting System Parameters (Network Bridging)	30
Checking MAC and Cipher Algorithms	31
Checking Password Authentication Settings	31
Installing the Required Operating System Packages	31
Removing Libraries	32
Configuring Time Synchronization	32
Configuring the Firewall	33
Configuring Proxy	34
Configuring DNS	34
Configuring the NFS Server	37
Disabling Swap Space	40
Creating Docker Thinpools	41
Enabling Installation Permissions for a sudo User	42
Installing CDF	44
Deploying ArcSight Fusion	46
Configuring the Cluster	46
Uploading Images to the Local Registry	48
Deploying Fusion	49
9 Deploying ArcSight Fusion in an Existing Cluster	51
Prerequisites	51
Deploying Fusion to an Existing Cluster	51
10 Post-Installation Configuration	53
Add Dashboard Roles and Permissions to the User Management Pod	53
Labeling Nodes	53
Labeling Nodes in a New Cluster	54
Labeling Nodes in an Existing Cluster	54
Connecting to an SMTP Server	55
Integrating Fusion Single Sign-On with an External SAML 2 Identity Provider	55
11 Verifying the Installation	59
Part III Configuring Fusion	61
12 Integrating Data and Users from Enterprise Security Manager	63
Understanding How ESM Users Access Fusion	63
Importing Users from ESM	63
Enabling Single Sign-on with ESM	64
Integrating Data from Enterprise Security Manager	65

13 Integrating Data from Intersect	67
14 Adding Users and Groups to the Dashboard	69
Part IV Managing ArcSight Fusion	71
15 Configuring the Dashboard	73
16 Using REST APIs with Fusion	75
17 Restarting Nodes in the Cluster	77
Restarting Nodes by Using Scripts	77
Restarting Nodes Manually	77
18 Resetting the CDF Administrator Password	79
19 Renewing CDF Certificates	81
Renewing Certificates Before Expiration	81
Renewing Certificates After Expiration	81
20 Creating and Adding CDF Certificate Authority	83

About This Book

This *Administrator Guide* provides information about deploying, configuring, and managing ArcSight Fusion.

Intended Audience

This book provides information for IT administrators who are responsible for managing the Fusion software and its environment. Usually, these individuals have experience in configuring servers and managing SIEM-related applications.

Additional Documentation

The ArcSight Fusion documentation library includes the following resources:

- ♦ *User Guide for ArcSight Fusion*, which is embedded in the product to provide both contextual Help and conceptual information
- ♦ *Release Notes for ArcSight Enterprise Security Manager*, which provides information about updates or new features available in the current release
- ♦ *Technical Requirements for ArcSight Fusion*, which provides information about the hardware and software requirements for installing Fusion

For the most recent version of this guide and other Fusion documentation resources, visit the [documentation for ArcSight Fusion](#) web page.

Contact Information

We want to hear your comments and suggestions about this book and the other documentation included with this product. You can use the **comment on this topic** link at the bottom of each page of the online documentation, or send an email to Documentation-Feedback@microfocus.com.

For specific product issues, contact Micro Focus Customer Care at <https://www.microfocus.com/support-and-services/>.

1 Understanding ArcSight Fusion

ArcSight Fusion enables you to visualize, identify, and analyze potential threats by incorporating intelligence from the multiple layers of security sources that might be installed in your security environment, such as:

- ♦ Real-time event monitoring and correlation with data from ArcSight Enterprise Security Manager (ESM)
- ♦ Analyzing end-user behavior with ArcSight Intersect

To help you get started, Fusion provides a Dashboard with a set of out-of-the-box widgets and dashboards. Users can organize the widgets into personalized dashboards.

Fusion scales to match the footprint of your environment. You can install Fusion [on the same server](#) as ESM, if there are enough spare resources on the server, or install [on multiple servers](#).


For information about using this product, see the *User Guide for ArcSight Fusion* embedded in the product or posted with the [documentation for ArcSight Fusion](#).

- ♦ [“Understanding the Fusion Architecture” on page 10](#)
- ♦ [“Understanding the Fusion Components” on page 10](#)

Understanding the Fusion Architecture

The Fusion environment incorporates several components that enable it to receive and display data from sources such as ESM. The following diagram helps you understand the software and components that comprise your Fusion environment.

Figure 1-1 Fusion Architecture



Understanding the Fusion Components

The following table describes the components incorporated into the Fusion environment.

Software/Component	Description
Data sources	Provide data to Fusion for display in the Dashboard
Network File System (NFS)	Stores some of the persistent data generated by Analytics and Fusion
Analytics (framework)	Provides user management, single sign-on, and other core services
Fusion (framework)	Provides dashboard configuration and management
Dashboard UI	Provides a browser-based interface for the Dashboard and User Management functions

Planning Installation and Deployment

This section helps you plan the installation and deployment of Fusion. Some components must be installed in a specific order because the process requires access to previously installed components.

- ♦ [Chapter 2, “Installation Checklist,” on page 13](#)
- ♦ [Chapter 3, “Understanding the Installation Process,” on page 15](#)
- ♦ [Chapter 4, “Installation Options,” on page 19](#)
- ♦ [Chapter 5, “Installing Fusion and ESM on the Same Node,” on page 21](#)

2 Installation Checklist

As part of your planning process, we recommend that you perform the tasks in this checklist in the order listed below.

	Checklist Items
<input type="checkbox"/>	<p>1. Review the description of the installation process.</p> <p>For more information, see Chapter 3, “Understanding the Installation Process,” on page 15.</p>
<input type="checkbox"/>	<p>2. Ensure that the computers, on which you are installing Fusion components, meet the specified hardware and software requirements.</p> <p>For more information, see the Technical Requirements for ArcSight Fusion.</p>
<input type="checkbox"/>	<p>3. Review the ports required for the installed and connected components.</p> <p>For more information, see “Ports Used” in the Technical Requirements for ArcSight Fusion.</p>
<input type="checkbox"/>	<p>4. Review the limitations and options associated with the installation.</p> <p>For more information, see Chapter 3, “Understanding the Installation Process,” on page 15.</p>
<input type="checkbox"/>	<p>5. (Conditional) If installing ESM and Fusion on the same node, add the ESM port to the firewall.</p> <p>For more information, see Chapter 5, “Installing Fusion and ESM on the Same Node,” on page 21.</p>
<input type="checkbox"/>	<p>6. Download the files for installing Fusion.</p> <p>For more information, see the Release Notes for ArcSight Enterprise Security Manager.</p>
<input type="checkbox"/>	<p>7. (Conditional) To collect data from ESM, install or upgrade your ESM environment.</p> <p>For more information, see the Installation Guide for ArcSight Enterprise Security Manager.</p>
<input type="checkbox"/>	<p>8. Install Fusion.</p> <p>For more information, see Chapter 7, “Installing ArcSight Fusion by Using Scripts,” on page 27.</p>
<input type="checkbox"/>	<p>9. Log in to Fusion to create the first administrative user.</p> <p>For more information, see Chapter 11, “Verifying the Installation,” on page 59.</p>
<input type="checkbox"/>	<p>10. Connect Fusion to your data sources:</p> <ul style="list-style-type: none">◆ Chapter 12, “Integrating Data and Users from Enterprise Security Manager,” on page 63◆ Chapter 13, “Integrating Data from Intersect,” on page 67

	Checklist Items
<input type="checkbox"/>	<p>11. Enable users to access Fusion:</p> <ul style="list-style-type: none">◆ Add or import users and groups◆ Create roles◆ Assign permissions <p>For more information, see the User Guide for ArcSight Fusion, which is also the context-sensitive Help in the product.</p>
<input type="checkbox"/>	<p>12. Create and share dashboards.</p> <p>For more information, see the User Guide for ArcSight Fusion, which is also the context-sensitive Help in the product.</p>

3 Understanding the Installation Process

Fusion is a container-based application. To deploy and manage Fusion, you must first install the **Container Deployment Foundation (CDF)**, which is a container and management module built on Kubernetes and Docker containers. You can deploy the [components](#) for Fusion either on a single-node or multi-node cluster setup.

If you already have an existing cluster for ArcSight Transformation Hub or ArcSight Investigate, you can deploy Fusion to the same cluster. Reusing existing clusters allows you to reduce costs and system management effort, when compared to deploying Fusion in a new cluster.

- ♦ [“Components Needed for Deployment” on page 15](#)
- ♦ [“Methods of Deployment” on page 16](#)
- ♦ [“Securing Communication Between Components” on page 16](#)

Components Needed for Deployment

The installation process deploys the following components:

- ♦ [“ArcSight Fusion” on page 15](#)
- ♦ [“Analytics Framework” on page 15](#)
- ♦ [“Documentation” on page 15](#)

ArcSight Fusion

Fusion incorporates intelligence from multiple layers of the security sources installed in your security environment and enables you to visualize, identify, and analyze potential threats that occur in your security environment.

Analytics Framework

Analytics Framework provides single sign-on and user management that enables you to search data stored in ArcSight Enterprise Security Manager and Interset from the Fusion user interface.

Documentation

The documentation component includes the [User Guide for ArcSight Fusion](#), which provides concepts, use cases, and contextual help for the users of ArcSight Fusion.

Methods of Deployment

You can deploy Fusion on one or more nodes, depending on the anticipated workload, and whether high availability deployment is required. The supported methods of deployment are:

- ♦ [“Single-Node Deployment” on page 16](#)
- ♦ [“Multi-Node Deployment” on page 16](#)
- ♦ [“High Availability Deployment” on page 16](#)

Single-Node Deployment

The single-node deployment method allows you to deploy all of the Fusion components on a single node. This method of deployment is suitable only for small workloads and where you do not need high availability. You can use the provided scripts, which [prepare your environment](#), [install CDF](#), and [deploy Fusion](#).

Multi-Node Deployment

For larger workloads, we recommend deploying Fusion in a multi-node cluster setup. The multi-node deployment method provides load balancing across several worker nodes and is scalable to handle large workloads. You can add multiple master nodes and worker nodes to scale. A multi-node deployment with three master and three worker nodes can be configured for HA support.

NOTE: Although, you can add worker nodes even after the installation, you can add master nodes only during the installation. Plan your deployment carefully before you start the installation process.

High Availability Deployment

For high availability deployment, you must set up a minimum of three master and three worker nodes, so that even in cases where two nodes are unavailable, the third node is still available. When only two master nodes are used, and the primary master node is offline for maintenance or upgrade, only a single master node remains available, which creates a single point of failure. If the available single master node fails, the cluster stops and cluster orchestration will not be possible until the master is back online.

NOTE: Although, you can add worker nodes even after the installation, you can add master nodes only during the installation. Plan your deployment carefully before you start the installation process.

Securing Communication Between Components

Determine the security mode (TLS/SSL) you want for establishing communication between the infrastructure components. The security mode must be the same across all components. Set up other Micro Focus components with the security mode you intend to use before connecting them.

NOTE: The secure communication described here applies only in the context of the components that relate to the Micro Focus container-based application you are using, which is specified in that application's documentation.

Changing the security mode after the deployment requires system downtime. If you do need to change the security mode after deployment, refer to the appropriate Administrator's Guide for the relevant component.

The following table lists Micro Focus components, preparations needed for secure communication with components, security modes, and where to find more information about the component.

Communication	Preparations needed...	Supported security modes
Analytics and Fusion to NFS Server	For optimal security, secure all NFS settings to allow only required hosts to connect to the NFS server.	
Web browser to NGINX (proxy)	No action is required.	TLS/SSL

Enabling FIPS in Nginx: No user action is required to enable FIPS for Nginx. The Nginx docker container is FIPS enabled by default. The FIPS-enabled Nginx server accepts TLS 1.2 connections using FIPS-compliant Cipher Suites.

4 Installation Options

You can install Fusion either by using the provided installation scripts or manually.

- ♦ [“Installation Using Scripts” on page 19](#)
- ♦ [“Manual Installation” on page 19](#)
- ♦ [“Deciding to Use the Scripts or Manual Installation Method” on page 19](#)

Installation Using Scripts

To enable an easier installation, Fusion provides scripts that automatically take care of all the prerequisites, software installations, and post-installation configurations. The scripts are applicable for single-node deployments where high availability is not needed. However, if you prefer to manually set the configurations and the installations because of your organization’s security policies, you can install Fusion manually in single-node deployments as well. The scripts configure the system to match the settings described for performing a manual installation.

The installation scripts require your environment to be in a specific state. Before deciding to use the installation scripts, review the [considerations for installation](#).

For information about installing Fusion by using scripts for a single-node deployment, see [Chapter 7, “Installing ArcSight Fusion by Using Scripts,” on page 27](#).

Manual Installation

In deployments with a larger workload where high availability is a requirement, you must manually perform all the necessary system configurations and software installations. However, you can use some of the installation scripts to make your tasks easier, then complete the rest of the configurations and installations manually.

For information about installing Fusion manually, see [Chapter 8, “Installing ArcSight Fusion Manually,” on page 29](#).

Deciding to Use the Scripts or Manual Installation Method

To determine whether to use the installation scripts or perform a manual installation, review the following considerations:

- ♦ The scripts install Fusion on the operating system with a default minimum installation. If you have any customizations on the operating system, we recommend that you perform the prerequisites manually and perform installation and post-installation configuration using scripts.

NOTE: When installing Fusion on the same machine as the ESM Manager, the installation scripts require that the ESM is installed on the system before running the installation scripts.

- ◆ The scripts install Fusion only on a singled-homed network (one that is connected with a single network link). If you have a dual-homed network (dual or redundant connections to a single Internet Service Provider), we recommend that you use the [manual installation](#) process.
- ◆ The scripts automatically tune the system for a single-node deployment with a small workload.
- ◆ The scripts install the cluster with a single master node and single worker node running on the same system. You can add worker nodes after the installation to scale and enable worker high availability.
- ◆ If you use the scripts, you cannot configure high availability for the master node. If you want high availability for the master node, we recommend that you use the [manual installation](#) process.
- ◆ The scripts disable the option to authorize Micro Focus to collect suite usage data.
- ◆ The scripts create NFS shares on the system used by the containers in the cluster. They configure the firewall to disable remote access to this NFS server. If you plan to add additional nodes to the cluster, you must enable remote access to the NFS server in the firewall.
- ◆ The scripts use the following paths by default:
 - ◆ To install Kubernetes: `/opt/kubernetes`
 - ◆ To create NFS shared directories: `/opt/arcsight/nfs/volumes`
- ◆ If you must use proxy in your environment, you must use the [manual installation](#) process.
- ◆ If your network is already using the subnets defined for the default CDF subnets, we recommend using the [manual installation](#) process. In this way, you can configure CDF to use a different subnet.

5 Installing Fusion and ESM on the Same Node

You can install Fusion on the same node as ESM, either manually or by using scripts.

- 1 Install ESM.

NOTE: Always install ESM before installing Fusion and CDF.

- 2 Add the ESM https port in iptables using the following commands:

- 2a To find your active zones, use the following command:

```
firewall-cmd --get-active-zones
```

- 2b To add the ESM port in iptables, use the following command:

```
firewall-cmd --zone=public --add-port=port_number/tcp --permanent
```

NOTE: This step enables you to access ESM externally (outside the firewall).

- 2c To reload the firewall so that the changes are applied, use the following command:

```
firewall-cmd --reload
```

- 3 Install Fusion.

For more information, see [Chapter 7, “Installing ArcSight Fusion by Using Scripts,”](#) on page 27 and [Chapter 8, “Installing ArcSight Fusion Manually,”](#) on page 29.

NOTE: While installing Fusion, specify a CDF API Server Port that does not use the same port as the ESM server (default 8443). For more information about Fusion ports, see the [Technical Requirements for ArcSight Fusion](#).



Installing CDF and Deploying Fusion

This section provides guidance for installing all components that are required to deploy Fusion in your environment.

- ♦ [Chapter 6, “Preparing Your Environment,” on page 25](#)
- ♦ [Chapter 7, “Installing ArcSight Fusion by Using Scripts,” on page 27](#)
- ♦ [Chapter 8, “Installing ArcSight Fusion Manually,” on page 29](#)
- ♦ [Chapter 9, “Deploying ArcSight Fusion in an Existing Cluster,” on page 51](#)
- ♦ [Chapter 10, “Post-Installation Configuration,” on page 53](#)
- ♦ [Chapter 11, “Verifying the Installation,” on page 59](#)

6 Preparing Your Environment

We provide the following script that prepares your environment to install all components on a single system:

```
./prepare-install-single-node-host.sh
```

For more information, see [Chapter 7, “Installing ArcSight Fusion by Using Scripts,”](#) on page 27.

7 Installing ArcSight Fusion by Using Scripts

You can use the installation scripts to install ArcSight Fusion on a single node. The installation scripts perform end-to-end installation starting from configuring prerequisites to completing post-installation configurations. For more information on running the scripts successfully, see [Chapter 8, “Installing ArcSight Fusion Manually,”](#) on page 29.

- ♦ [“Prerequisites”](#) on page 27
- ♦ [“Understanding the Installation Scripts”](#) on page 27
- ♦ [“Using the Installation Scripts”](#) on page 27

Prerequisites

Ensure that your environment meets the hardware and software requirements described in the [Technical Requirements for ArcSight Fusion](#).

Understanding the Installation Scripts

The installation scripts automatically take care of all the prerequisites, software installations, and post-installation configurations.

Script	Purpose
<code>./prepare-install-single-node-host.sh</code>	Installs all the necessary packages and configures the prerequisites
<code>./install-single-node.sh</code>	Installs ArcSight Fusion
<code>./install-single-node-post.sh</code>	Performs post-installation configurations, such as labeling the nodes

Using the Installation Scripts

IMPORTANT: The information in this section applies only when your deployment does not need high availability.

To install Fusion by using the scripts:

- 1 Log in to the master node as `root`.
- 2 Change to the directory where you downloaded the ArcSight Fusion installation files.

Example: `cd /opt`

For information about downloading the ArcSight Fusion installation files, see “Downloading Fusion” in the [Release Notes for ArcSight Enterprise Security Manager](#).

3 Extract the downloaded ArcSight Fusion installation files using the following command:

```
tar zxvf fusion-installer-1.0.0.nn.tar.gz
```

4 Execute the scripts in the following order:

4a `./prepare-install-single-node-host.sh`

4b `./install-single-node.sh`

4c `./install-single-node-post.sh`

8

Installing ArcSight Fusion Manually

This chapter provides information about manually installing Fusion and the required software.

- ♦ [“Preparing Your Environment for CDF” on page 29](#)
- ♦ [“Installing CDF” on page 44](#)
- ♦ [“Deploying ArcSight Fusion” on page 46](#)

Preparing Your Environment for CDF

The procedures in this section enable you to configure your environment for a successful installation of the Container Deployment Foundation (CDF).

- ♦ [“Configuring the Nodes” on page 29](#)
- ♦ [“Setting System Parameters \(Network Bridging\)” on page 30](#)
- ♦ [“Checking MAC and Cipher Algorithms” on page 31](#)
- ♦ [“Checking Password Authentication Settings” on page 31](#)
- ♦ [“Installing the Required Operating System Packages” on page 31](#)
- ♦ [“Removing Libraries” on page 32](#)
- ♦ [“Configuring Time Synchronization” on page 32](#)
- ♦ [“Configuring the Firewall” on page 33](#)
- ♦ [“Configuring Proxy” on page 34](#)
- ♦ [“Configuring DNS” on page 34](#)
- ♦ [“Configuring the NFS Server” on page 37](#)
- ♦ [“Disabling Swap Space” on page 40](#)
- ♦ [“Creating Docker Thinpools” on page 41](#)
- ♦ [“Enabling Installation Permissions for a sudo User” on page 42](#)

Configuring the Nodes

For multi-node deployment, consider the following when configuring master and worker nodes:

- ♦ Deploy master and worker nodes on virtual machines. Since most of the processing occurs on worker nodes, we recommend that you deploy worker nodes on physical servers.
- ♦ Keep the host system configuration identical across master and worker nodes.
- ♦ When using virtual machines, ensure that:
 - ♦ Resources are reserved and not shared
 - ♦ UUID and MAC addresses are static because dynamic addresses cause the Kubernetes cluster to fail

- ◆ Install all master and worker nodes in the same subnet.
- ◆ Add more worker nodes rather than installing bigger and faster hardware.

Using more worker nodes enables you to perform maintenance on your cluster nodes with minimal impact to uptime. Adding more nodes also helps with predicting costs due to new hardware.

For high availability, consider the following when configuring master and worker nodes:

- ◆ Create a virtual IP that is shared by all master nodes and ensure that virtual IP is under the same subnet. The VIP must not respond when pinged before you install Fusion.
- ◆ Install all master and worker nodes in the same subnet.

Setting System Parameters (Network Bridging)

Ensure that the `br_netfilter` module is installed on all master and worker nodes before changing system settings.

You can either run the following scripts that set system parameters automatically or you can set the system parameters manually:

- ◆ `/opt/<ArcSight_Fusion_Installer>/scripts/prereq_sysctl_conf.sh`
- ◆ `/opt/<ArcSight_Fusion_Installer>/scripts/prereq_rc_local.sh`

Perform the following steps on all the master and worker nodes to set the system parameters manually:

- 1 Log in to the node.
- 2 Check whether the `br_netfilter` module is enabled:


```
lsmod |grep br_netfilter
```
- 3 If there is no return value and the `br_netfilter` module is not installed, install it:


```
modprobe br_netfilter
echo "br_netfilter" > /etc/modules-load.d/br_netfilter.conf
```
- 4 Open the `/etc/sysctl.conf` file.
- 5 Ensure that the following system parameters are set:


```
net.bridge.bridge-nf-call-iptables=1
net.bridge.bridge-nf-call-ip6tables=1
net.ipv4.ip_forward = 1
net.ipv4.tcp_tw_recycle = 0
kernel.sem=50100 128256000 50100 2560
```
- 6 Save the `/etc/sysctl.conf` file.
- 7 Apply the updates to the node:


```
/sbin/sysctl -p
```

Checking MAC and Cipher Algorithms

To configure MAC and Cipher algorithms manually, ensure that the `/etc/ssh/sshd_config` files on every master and worker node are configured with at least one of the following values, which lists all supported algorithms. Add only the algorithms that meet the security policy of your organization.

- ♦ For MAC algorithms: `hmac-sha1`, `hmac-sha2-256`, `hmac-sha2-512`, `hmac-sha1-96`
- ♦ For Cipher algorithms: `3des-cbc`, `aes128-cbc`, `aes192-cbc`, `aes256-cbc`, `aes128-ctr`, `aes192-ctr`, `aes256-ctr`, `arcfour128`, `arcfour256`, `blowfish-cbc`

For example, you could add the following lines to the `/etc/ssh/sshd_config` file on all master and worker nodes:

```
MACs hmac-sha2-256,hmac-sha2-512
```

```
Ciphers aes128-cbc,aes192-cbc,aes256-cbc,aes128-ctr,aes192-ctr,aes256-ctr
```

Checking Password Authentication Settings

If you plan to use a user name and password authentication for adding cluster nodes during the installation, ensure that the `PasswordAuthentication` parameter in the `/etc/ssh/sshd_config` file is set to `yes`. There is no need to check the password authentication setting when you add cluster nodes using a user name and key authentication.

To ensure that the password authentication is enabled, perform the following steps on every master and worker node:

- 1 Log in to the cluster node.
- 2 Open the `/etc/ssh/sshd_config` file.
- 3 Check whether the `PasswordAuthentication` parameter is set to `yes`. If not, set the parameter to `yes` as follows:

```
PasswordAuthentication yes
```

- 4 Restart the `sshd` service:

```
systemctl restart sshd.service
```

Installing the Required Operating System Packages

Ensure that the packages listed in the following table are installed on appropriate nodes. These packages are available in the standard `yum` repository.

Package	Nodes
<code>device-mapper-libs</code>	Master and worker
<code>java-1.8.0-openjdk</code>	Master
<code>libgcrypt</code>	Master and worker
<code>libseccomp</code>	Master and worker
<code>libtool-ltdl</code>	Master and worker

Package	Nodes
net-tools	Master and worker
nfs-utils	Master and worker
rpcbind	Master node, worker node, and NFS server
systemd-libs (version >= 219)	Master and worker
unzip	Master and worker
httpd-tools	Master and worker
conntrack-tools	Master and worker
lvm2	Master and worker
curl	Master and worker
libtool-libs	Master and worker
openssl	Master and worker
socat	Master and worker
container-selinux	Master and worker

You can either run the `/opt/<ArcSight_Fusion_Installer>/scripts/prereq_1_required_packages.sh` script that installs the required OS packages automatically or install the required OS packages manually.

To install the packages manually:

- 1 Log in to the master or worker nodes.
- 2 Verify whether the package exists:

```
yum list installed <package name>
```
- 3 (Conditional) If the package is not installed, install the required package:

```
yum -y install <package name>
```

Removing Libraries

Remove libraries that prevent Ingress from starting and confirm the removal when prompted:

```
yum remove rsh rsh-server vsftpd
```

Configuring Time Synchronization

You must implement a Network Time Protocol (NTP) to synchronize time on all nodes in the cluster. To implement this protocol, use chrony. Ensure that chrony is running on all nodes in the cluster. By default chrony is installed on some versions of RHEL.

You can either run the `/opt/<ArcSight_Fusion_Installer>/scripts/prereq_synchronize_time.sh` script that synchronizes time automatically or configure the time synchronization manually.

To configure the time synchronization manually:

- 1 Verify chrony configuration:

```
chronyc tracking
```

- 2 (Conditional) If chrony is not installed, install chrony:

```
yum install chrony
```

- 3 Start and enable chrony:

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

- 4 Synchronize the operating system time with the NTP server:

```
chronyc makestep
```

- 5 Restart the chronyd daemon:

```
systemctl restart chronyd
```

- 6 Check the server time synchronization:

```
timedatectl
```

- 7 Synchronize the hardware time:

```
hwclock -w
```

Configuring the Firewall

Ensure that the `firewalld.service` is enabled and running on all nodes.

You can either run the `/opt/<ArcSight_Fusion_Installer>/scripts/prereq_firewall.sh` script that configures the firewall automatically or configure the firewall manually.

When the firewall is enabled, you must also enable the masquerade settings. To enable masquerade settings:

- 1 Check whether the masquerade setting is already enabled:

```
firewall-cmd --query-masquerade
```

If the command returns `yes`, then masquerade is enabled.

If the command returns `no`, then masquerade is disabled.

- 2 (Conditional) If the masquerade setting is not enabled, enable masquerade:

```
firewall-cmd --add-masquerade --permanent  
firewall-cmd --reload
```

Configuring Proxy

Ensure that the cluster should have no access to the Internet and that the proxy settings (`http_proxy`, `https_proxy`, and `no_proxy`) are not set. However, if a connection with the Internet is needed and you already specified a proxy server for http and https connection, you must correctly configure `no_proxy`.

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

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

For example:

- ♦

```
export http_proxy="http://web-proxy.example.net:8080"
export https_proxy="http://web-proxy.example.net:8080"
export
no_proxy="localhost,127.0.0.1,node1.swinfra.net,10.94.235.231,node2.sw
infra.net,10.94.235.232,node3.swinfra.net,10.94.235.233,node3.swinfra.
net,10.94.235.233,node4.swinfra.net,10.94.235.234,node5.swinfra.net,10
.94.235.235,node6.swinfra.net,10.94.235.236,ha.swinfra.net
10.94.235.200"
```
- ♦

```
export http_proxy="http://web-proxy.eu.example.net:8080"
export
https_proxy="localhost,127.0.0.1,swinfra.net,10.94.235.231,10.94.235.2
32,10.94.235.233,10.94.235.233,10.94.235.234,10.94.235.235,10.94.235.2
36,10.94.235.200"
```

NOTE: Incorrect configuration of proxy settings has proven to be a frequent installation problem. To verify that proxy settings are configured properly on all master and worker nodes, run the following command and ensure that the output corresponds to the recommendations:

```
echo $http_proxy, $https_proxy, $no_proxy
```

If the firewall is turned off, the installation process will generate a warning. To prevent the warning, set the CDF install parameter `--auto-configure-firewall` to `true`.

Configuring DNS

Ensure that the host name resolution through Domain Name System (DNS) is working across all nodes in the cluster, including correct forward and reverse DNS lookups. Host name resolution must not be performed through `/etc/hosts` file settings.

You can either run the `<download_directory>/scripts/prereq_disable_ipv6.sh` script that configures DNS automatically or configure DNS manually.

Ensure that all nodes are configured with a Fully Qualified Domain Name (FQDN) and are in the same subnet. Transformation Hub uses the host system FQDN as its Kafka `advertised.host.name`. If the FQDN resolves successfully in the Network Address Translation (NAT) environment, producers and consumers will function correctly. If there are network-specific issues resolving FQDN through NAT, DNS will need to be updated to resolve these issues.

- ◆ Transformation Hub supports ingestion of event data that contains both IPv4 and IPv6 addresses. However, its infrastructure cannot be installed in an IPv6-only network.
- ◆ `localhost` must not resolve to an IPv6 address, such as, `::1` – this is the default state. The installation process expects only IPv4 resolution to IP address `127.0.0.1`. Comment out any `::1` reference.
- ◆ The initial master node host name must not resolve to multiple IPv4 addresses and this includes lookup in `/etc/hosts`.
- ◆ [“Testing Forward and Reverse DNS Lookup” on page 35](#)
- ◆ [“Understanding Kubernetes Network Subnet Settings” on page 36](#)

Testing Forward and Reverse DNS Lookup

Test that the forward and reverse lookup records for all servers were properly configured.

To test the forward lookup, run the following commands on every master and worker node in the cluster and on every producer and consumer host system, including:

- ◆ All master nodes: `master1.yourcompany.com, ..., mastern.yourcompany.com`
- ◆ All worker nodes: `worker1.yourcompany.com, ..., workern.yourcompany.com`
- ◆ Your ArcMC nodes: `arcmc1.yourcompany.com, ..., arcmcn.yourcompany.com`

Use the `nslookup` or `host` commands to verify your DNS configuration.

NOTE: Do not use the `ping` command.

You must run the `nslookup` commands on every server specified in your `/etc/resolv.conf` file. Every server must be able to perform forward and reverse lookup properly and return identical results.

If you have a public DNS server specified in your `/etc/resolv.conf` file, such as the Google public DNS server `8.8.8.8` or `8.8.4.4`, you must remove this server from your DNS configuration.

Run the commands as follows. Expected sample output is shown below each command.

- ◆ `hostname`
`master1`
- ◆ `hostname -s`
`master1`
- ◆ `hostname -f`
`master1.yourcompany.com`

- ◆ `hostname -d`
`yourcompany.com`
- ◆ `nslookup master1.yourcompany.com`
Server: 192.168.0.53
Address: 192.168.0.53#53
Address: 192.168.0.1
Name: master1.example.com
- ◆ `nslookup master1`
Server: 192.168.0.53
Address: 192.168.0.53#53
Name: master1.example.com
Address: 192.168.0.1
- ◆ `nslookup 192.168.0.1`
Server: 192.168.0.53
Address: 192.168.0.53#53
1.0.168.192.in-addr.arpa name = master1.example.com.

Understanding Kubernetes Network Subnet Settings

The Kubernetes network subnet is controlled by the `--POD_CIDR` and `-SERVICE_CIDR` parameters to the Container Deployment Foundation (CDF) installation portal.

The `--POD_CIDR` parameter specifies the network address range for Kubernetes pods. The address range specified in the `--POD_CIDR` parameter must not overlap with the IP range assigned for Kubernetes services, which is specified in the `-SERVICE_CIDR` parameter. The expected value is a Classless Inter-Domain Routing (CIDR) format IP address. CIDR notation comprises an IP address, a slash (/) character, and a network prefix (a decimal number). The minimum useful network prefix is /24 and the maximum useful network prefix is /8. The default value is `172.16.0.0/16`.

For example:

```
POD_CIDR=172.16.0.0/16
```

The `CIDR_SUBNETLEN` parameter specifies the size of the subnet allocated to each host for Kubernetes pod network addresses. The default value is dependent on the value of the `POD_CIDR` parameter, as described in the following table.

POD_CIDR Prefix	POD_CIDR_SUBNETLEN defaults	POD_CIDR_SUBNETLEN allowed values
/8 to /21	/24	/(POD_CIDR prefix + 3) to /27
/22 to /24	/(POD_CIDR prefix + 3)	/(POD_CIDR prefix + 3) to /27

Smaller prefix values indicate a larger number of available addresses. The minimum useful network prefix is /27 and the maximum useful network prefix is /12. The default value is `172.17.17.0/24`.

Change the default `POD_CIDR` or `CIDR_SUBNETLEN` values only when your network configuration requires you to do so. You must also ensure that you have sufficient understanding of the flannel network fabric configuration requirements before you make any changes.

Configuring the NFS Server

Container Deployment Foundation (CDF) requires an NFS server to maintain state information about the infrastructure and to store other pertinent data.

For high availability, NFS must run on a highly available external server in the case of a dedicated master deployment having a minimum of three master nodes. For optimal security, secure all NFS settings to allow only required hosts to connect to the NFS server.

- ♦ [“Prerequisites” on page 37](#)
- ♦ [“Creating NFS Shared Directories” on page 38](#)
- ♦ [“Exporting the NFS Configuration” on page 39](#)
- ♦ [“Verifying NFS Configuration” on page 39](#)
- ♦ [“Setting Up NFS By Using the Script” on page 40](#)

Prerequisites

The prerequisites for configuring the NFS server are listed below:

- ♦ Ensure that the ports 111, 2049, and 20048 are open on the NFS server for communication.
- ♦ Enable the `rpcbind` and `nfs-server` package by executing the following commands on your NFS server:

```
systemctl enable rpcbind
systemctl start rpcbind
systemctl enable nfs-server
systemctl start nfs-server
```

- ♦ Create and configure the following shared directories.

Directory	Description
<code><NFS_ROOT_DIRECTORY>/itom-vol</code>	This is the CDF NFS root folder, which contains the CDF database and files. The disk usage will grow gradually.
<code><NFS_ROOT_DIRECTORY>/db-single-vol</code>	This volume is available only if you did not choose PostgreSQL High Availability (HA) for the CDF database setting. It is for the CDF database. During the installation you will not choose the PostgreSQL database HA option.
<code><NFS_ROOT_DIRECTORY>/db-backup-vol</code>	This volume is used for backup and restoration of the CDF PostgreSQL database. Its sizing is dependent on the implementation's processing requirements and data volumes.
<code><NFS_ROOT_DIRECTORY>/itom-logging-vol</code>	This volume stores the log output files of CDF components. The required size depends on how long the log will be kept.
<code><NFS_ROOT_DIRECTORY>/arcsight-vol</code>	This volume stores the component installation packages.

Creating NFS Shared Directories

To create NFS shared directories:

- 1 Log in to the NFS server as `root`.
- 2 Create the following:
 - ♦ **Group:** `arcsight` with a GID 1999
 - ♦ **User:** `arcsight` with a UID 1999
 - ♦ **NFS root directory:** Root directory under which you can create all NFS shared directories.

Example (NFS_Root_Directory): `/opt/arcsight/nfs/volumes`

- 3 (Conditional) If you have previously installed any version of CDF, you must remove all NFS directories using the following command for each directory:

```
rm -rf <path to NFS directory>
```

Example:

```
rm -rf /opt/arcsight/nfs/volumes/itom-vol
```

- 4 Create each NFS shared directory using the command:

```
mkdir -p <path to NFS directory>
```

Example:

```
mkdir -p /opt/arcsight/nfs/volumes/itom-vol
```

- 5 For each NFS directory, set the permission to 755 using the command:

```
chmod -R 755 <path to NFS directory>
```

Example:

```
chmod -R 755 /opt/arcsight/nfs/volumes/itom-vol
```

- 6 For each NFS directory, set the ownership to UID 1999 and GID 1999 using the command:

```
chown -R 1999:1999 <path to NFS directory>
```

Example:

```
chown -R 1999:1999 /opt/arcsight/nfs/volumes/itom-vol
```

If you use a UID/GID other than 1999/1999, provide it during the CDF installation in the installation script arguments `--system-group-id` and `--system-user-id`.

Exporting the NFS Configuration

For every NFS volume, run the following set of commands on the External NFS server based on the IP address. You will need to export the NFS configuration with the appropriate IP address for the NFS mount to work properly.

- 1 For every node in the cluster, you must update the configuration to grant the node access to the NFS volume shares.

For example:

```
/opt/arcsight/nfs/volumes/arcsight-vol 192.168.1.0/  
24(rw, sync, anonuid=1999, anongid=1999, all_squash)  
  
/opt/arcsight/nfs/volumes/itom-vol 192.168.1.0/  
24(rw, sync, anonuid=1999, anongid=1999, all_squash)  
  
/opt/arcsight/nfs/volumes/db-single-vol 192.168.1.0/  
24(rw, sync, anonuid=1999, anongid=1999, all_squash)  
  
/opt/arcsight/nfs/volumes/itom-logging-vol 192.168.1.0/  
24(rw, sync, anonuid=1999, anongid=1999, all_squash)  
  
/opt/arcsight/nfs/volumes/db-backup-vol 192.168.1.0/  
24(rw, sync, anonuid=1999, anongid=1999, all_squash)
```

- 2 Modify the `/etc/exports` file and run the following command:

```
exportfs -ra
```

If you add more NFS shared directories later, you must restart the NFS service.

Verifying NFS Configuration

To verify NFS configuration:

- 1 Create the NFS directory under `/mnt`.
- 2 Mount the NFS directory on your local system by using the command:
 - ♦ **NFS v3:** `mount -t nfs 192.168.1.25:/opt/arcsight/nfs/volumes/arcsight-vol /mnt/nfs`
 - ♦ **NFS v4:** `mount -t nfs4 192.168.1.25:/opt/arcsight/nfs/volumes/arcsight-vol /mnt/nfs`
- 3 After creating all the directories, run the following commands on the NFS server:

```
exportfs -ra
systemctl restart rpcbind
systemctl enable rpcbind
systemctl restart nfs-server
systemctl enable nfs-server
```

Setting Up NFS By Using the Script

IMPORTANT: The information in this section applies only for non-high-availability and single-node deployments.

You can either run the `/opt/<ArcSight_Fusion_Installer>/scripts/preinstall_create_nfs_share.sh` script that sets up the NFS automatically or set up NFS manually.

To set up NFS manually:

- 1 Copy `setupNFS.sh` to the NFS server.

The `setupNFS.sh` file is located on the master node in the `<download_directory>/fusion-installer-x.x.x.x/installers/cdf-x.x.x.x/scripts` folder.

- 2 (Conditional) If you are using the default UID/GID, use the command:

```
sh setupNFS.sh <path_to_nfs_directory>/volumes/volume_name
```

- 3 (Conditional) If you are using a non-default UID/GID, use the command:

```
sh setupNFS.sh <path_to_nfs_directory>/volumes/volume_name true <uid>
<gid>
```

- 4 Restart the NFS service:

```
systemctl restart nfs
```

Disabling Swap Space

You must disable swap space on all master and worker nodes, excluding the node that has Vertica.

- 1 Log in to the node where you want to disable swap space.

- 2 Run the following command:

```
swapoff -a
```

- 3 In the `/etc/fstab` file, comment out the lines that contain `swap` as the disk type and save the file.

For example:

```
#/dev/mapper/centos_shcentos72x64-swap swap
```


Creating Docker Thinpools

Optionally, to improve performance of Docker processing, set up a thinpool on each master and worker node. Before setting up a thinpool on each node, create a single disk partition on the node, as explained below.

For the thinpool device for Docker (for example, **sdb1**) the minimum physical volume size is 30 GB.

- ♦ [“Creating a New Partition” on page 41](#)
- ♦ [“Setting Up a Thinpool for Docker” on page 41](#)

Creating a New Partition

- 1 Log in to the node.
- 2 Run the command:

```
fdisk <name of the new disk device that was added>
```

Example:

```
# fdisk /dev/sdb1
```
- 3 Enter **n** to create a new partition.
- 4 When prompted, enter the partition number, sector, type (**Linux LVM**), and size for the first partition. To select the Linux LVM partition type:
 - ♦ Enter **t** to change the default partition type to Linux LVM
 - ♦ Type **L** to list the supported partition types
 - ♦ Type **8e** to select the Linux LVM type
- 5 When prompted, enter the partition number, sector, type (**Linux LVM**), and size for the second partition.
- 6 Type **p** to view the partition table.
- 7 Type **w** to save the partition table to disk.
- 8 Type **partprobe**.

Setting Up a Thinpool for Docker

- 1 Create a physical volume with the following command:

```
# pvcreate [physical device name]
```

Example:

```
# pvcreate /dev/sdb1
```
- 2 Create a volume group with the following command:

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

Example:

```
# vgcreate docker /dev/sdb1
```
- 3 Create a logical volume (LV) for the thinpool and bootstrap with the following command:

```
# lvcreate [logical volume name] [volume group name]
```

For example, the data LV is 95% of the 'Docker' volume group size. (Leaving free space allows for automatic expanding of either the data or metadata if space is running low, as a temporary measure.)

```
# lvcreate --wipesignatures y -n thinpool docker -l 95%VG
# lvcreate --wipesignatures y -n thinpoolmeta docker -l 1%VG
```

4 Convert the pool to a thinpool with the following command:

```
# lvconvert -y --zero n -c 512K --thinpool docker/thinpool --
poolmetadata docker/thinpoolmeta
```

Optionally, you can configure the auto-extension of thinpools using an lvm profile.

4a Open the lvm profile.

4b Specify a value for the parameters `thin_pool_autoextend_threshold` and `thin_pool_autoextend_percent`, each of which represents a percentage of the space used.

For example:

```
activation {
  thin_pool_autoextend_threshold=80
  thin_pool_autoextend_percent=20
}
```

4c Apply the lvm profile with the following command:

```
# lvchange --metadataprofile docker-thinpool docker/thinpool
```

4d Verify that the lvm profile is monitored with the following command:

```
# lvs -o+seg_monitor
```

4e Clear the graph driver directory with the following command, if Docker was previously started:

```
# rm -rf /var/lib/docker/*
```

4f Monitor the thinpool and volume group free space with the following commands:

```
# lvs
# lvs -a
# vgs
```

4g Check the logs to see the auto-extension of the thinpool when it hits the threshold:

```
# journalctl -fu dm-event.service
```

Enabling Installation Permissions for a sudo User

If you choose to install CDF as a `sudo` user, the root user must grant non-root (`sudo`) users installation permission before they can perform the installation. Ensure that the provided user has permission to execute scripts under temporary directory `/tmp` on all master and worker nodes.

There are two distinct file edits that need to be performed: First on the initial master node only, and then on all remaining master and worker nodes.

- ♦ [“Edit the sudoers File on the Initial Master Node”](#) on page 43
- ♦ [“Edit the sudoers File on the Remaining Master and Worker Nodes”](#) on page 44

Edit the `sudoers` File on the Initial Master Node

Make the following modifications only on the initial master node.

IMPORTANT: In the following commands you must ensure that there no more than a single space character after each comma that delimits parameters. Otherwise, you may get an error similar to this when you attempt to save the file:

```
>>> /etc/sudoers: syntax error near line nn<<<
```

1 Log in to the initial master node as the `root` user.

2 Open the `/etc/sudoers` file using `Visudo`.

3 Add the following `Cmdn_Alias` line to the command aliases group in the `sudoers` file:

```
Cmdn_Alias CDFINSTALL = <CDF_installation_package_directory>/scripts/  
precheck.sh, <CDF_installation_package_directory>/install, <K8S_HOME>/  
uninstall.sh, /usr/bin/kubect1, /usr/bin/docker, /usr/bin/mkdir, /bin/  
rm, /bin/su, /bin/chmod, /bin/tar, <K8S_HOME>/scripts/uploadimages.sh, /  
bin/chown
```

3a Replace `<CDF_installation_package_directory>` with the directory where you unzipped the installation package.

For example: `/tmp/cdf-2019.05.0xxx`.

3b Replace `<K8S_HOME>` with the value defined from a command line. By default, `<K8S_HOME>` is `/opt/arcsight/kubernetes`.

4 Add the following lines to the `wheel` users group, replacing `<username>` with your `sudo` user password:

```
%wheel ALL=(ALL) ALL  
cdfuser ALL=NOPASSWD: CDFINSTALL
```

Example:

```
Defaults: root !requiretty
```

5 Locate the `secure_path` line in the `sudoers` file and ensure that the following paths are present:

```
Defaults secure_path = /sbin:/bin:/usr/sbin:/usr/bin
```

By doing this, the `sudo` user can execute the `showmount`, `curl`, `ifconfig`, and `unzip` commands when installing CDF.

6 Save the file.

Installing Components Using the `sudo` User

After completing the modifications to the `sudoers` file as described above, perform the following steps:

- 1 Log in to the initial master node as the non-root `sudo` user to perform the installation.
- 2 Download the installation files to a directory where the non-root `sudo` user has write permissions.
- 3 Run CDF using the `sudo` command.

Edit the `sudoers` File on the Remaining Master and Worker Nodes

Make the following modifications only on the remaining master and worker nodes.

IMPORTANT: In the following commands you must ensure that there is, at most, a single space character after each comma that delimits parameters. Otherwise, you may get an error similar to this when you attempt to save the file.

```
>>> /etc/sudoers: syntax error near line nn<<<
```

- 1 Log in to each master and worker node.
- 2 Open the `/etc/sudoers` file.
- 3 Add the following `Cmnd_Alias` line to the command aliases group in the `sudoers` file.

```
Cmnd_Alias CDFINSTALL = /tmp/scripts/pre-check.sh,  
<ITOM_Suite_Foundation_Node>/install, <K8S_HOME>/uninstall.sh, /usr/  
bin/kubect1, /usr/bin/docker,/usr/bin/mkdir, /bin/rm, /bin/su, /bin/  
chmod, /bin/tar, <K8S_HOME>/scripts/uploadimages.sh, /bin/chown
```

- 3a Replace `<ITOM_Suite_Foundation_Node>` with the directory where you unzipped the installation package.

For example: `/tmp/ITOM_Suite_Foundation_2019.05.0xxx`

- 3b Replace `<K8S_HOME>` with the value defined from a command line. By default, `<K8S_HOME>` is `/opt/arcsight/kubernetes`.

- 4 Add the following lines to the `wheel` users group, replacing `<username>` with your `sudo` user password:

```
%wheel ALL=(ALL) ALL  
cdfuser ALL=NOPASSWD: CDFINSTALL
```

Example:

```
Defaults: root !requiretty
```

- 5 Locate the `secure_path` line in the `sudoers` file and ensure that the following paths are present:

```
Defaults secure_path = /sbin:/bin:/usr/sbin:/usr/bin
```

By doing this, the `sudo` user can execute the `showmount`, `curl`, `ifconfig`, and `unzip` commands when installing CDF.

- 6 Save the file.

Repeat the process for each remaining master and worker node.

Installing CDF

This section provides guidance for installing the Container Deployment Foundation (CDF).

NOTE: You can install CDF as a root user or `sudo` user. However, if you choose to install as a `sudo` user, you must first configure installation permissions from the root user.

- 1 Log in to the master node as the `root` or `sudo` user.
- 2 Change to the directory where you downloaded the installation files. For information about downloading the installation files, see [Part II, “Installing CDF and Deploying Fusion,” on page 23](#).

```
cd <download_directory>/fusion-installer-x.x.x.x/installers/cdf-  
<year>.<month>.<build_number>-<version>
```

For example:

```
cd opt/fusion-installer-x.x.x.x/installers/cdf-2020.02.00120-2.2.0.2
```

- 3 (Conditional) If you are installing CDF on a node where ESM Manager is not installed, use the command:

```
./install -m <metadata_file_path>  
--k8s-home <cdf_installer_directory>  
--nfs-server <NFS_server_IP_address>  
--nfs-folder <NFS_ITOM_volume_file_path>  
--registry-orgname srg
```

Example:

```
./install -m /opt/fusion-installer-x.x.x.x/suite_images/arcsight-  
installer-metadata-x.x.x.x.tar  
--k8s-home /opt/arcsight/kubernetes  
--nfs-server <NFS_server_IP_address>  
--nfs-folder /opt/arcsight/nfs/volumes/itom-vol  
--registry-orgname srg
```

- 4 (Conditional) If you are installing CDF on the same machine as ESM Manager, complete the following action:

If ESM is installed with the default port (8443), add the following command argument to configure the CDF API server port (default 8443) to use a different port:

```
--master-api-ssl-port 7443
```

Example:

```
./install -m /opt/fusion-installer-x.x.x.x/suite_images/arcsight-  
installer-metadata-x.x.x.x.tar  
--k8s-home /opt/arcsight/kubernetes  
--nfs-server <NFS_server_IP_address>  
--master-api-ssl-port 7443  
--nfs-folder /opt/arcsight/nfs/volumes/itom/itom_vol  
--registry-orgname srg
```

- 5 (Conditional) For high availability, use the following command:

```
./install -m <metadata_file_path>  
--k8s-home <NFS_server_IP_address>  
--nfs-folder <NFS_ITOM_volume_file_path>  
--registry-orgname srg  
--ha-virtual-ip <HA_virtual_IP_address>
```

Example:

```
./install -m /opt/fusion-installer-x.x.x.x/suite_images/arc sight-  
installer-metadata-x.x.x.x.tar  
--k8s-home /opt/arc sight/kubernetes  
--nfs-server <NFS_server_IP_address>  
--nfs-folder /opt/arc sight/nfs/volumes/itom-vol  
--registry-orgname srg  
--ha-virtual-ip <NFS_server_IP_address>
```

- 6 When prompted, specify the administrator password. This password is required to log in to the CDF Management Portal as an administrator.

The CDF Management Portal enables you to deploy Fusion and all required software in a cluster.

- 7 Change the default CA certificate that is generated during the installation. For steps to change the CA certificate, see [Changing the CA of CDF](#).

A self-signed CA certificate is generated during the installation of CDF by default. After deploying Fusion, the pods of the deployed products use the certificates generated by the CA on pod startup. If you change the CA certificate after deployment, you will have to uninstall and reinstall all the software. So, we recommend changing the CA certificate before deploying Fusion.

- 8 To deploy Fusion and all required software, continue with the section [“Deploying ArcSight Fusion”](#) on page 46.

Deploying ArcSight Fusion

This section provides information about using the CDF Management Portal to deploy ArcSight Fusion.

- ♦ [“Configuring the Cluster”](#) on page 46
- ♦ [“Uploading Images to the Local Registry”](#) on page 48
- ♦ [“Deploying Fusion”](#) on page 49

Configuring the Cluster

- 1 Open a new tab in a supported web browser.
- 2 Specify the URL for the CDF Management Portal:

```
https://<Fusion-server>:3000
```

NOTE: Use port *3000* when you are setting up the CDF for the first time. After the initial setup, use port *5443* to access the CDF Management Portal.

Use the fully qualified domain name of the host that you specified in the **Connection** step during the CDF configuration. Usually, this is the master node’s FQDN.

- 3 Log in to the CDF Management Portal with the credentials of the administrative user that you provided during installation.
- 4 Select the metadata file version in **version** and click **Next**.
- 5 Read the license agreement and select **I agree**.
- 6 Click **Next**.

- 7 On the Capabilities page, select the following and click **Next**:
 - ◆ ArcSight Fusion
 - ◆ Analytics
- 8 On the Database page, retain the default values and click **Next**.
- 9 On the Deployment Size page, select the required cluster and click **Next**.
 - 9a (Conditional) For worker node configuration, select **Medium Cluster**.
- 10 On the Connection page, an external host name is automatically populated. This is resolved from the virtual IP (VIP) specified during the CDF installation (--ha-virtual-ip parameter). Confirm that the VIP is correct and then click **Next**.
- 11 (Conditional) If you want to set up high availability, select **Make master highly available** and add at least two additional master nodes on the Master High Availability page.

IMPORTANT: If you do not configure high availability in this step, you cannot add master nodes and configure high availability after installation.

On the Add Master Node page, specify the following details:

- ◆ **Host:** Fully qualified domain name (FQDN) of the node you are adding.
- ◆ **Ignore Warnings:** If selected, the CDF Management Portal will ignore any warnings that occur during the pre-checks on the server. If deselected, the add node process will stop and a window will display any warning messages. We recommend that you start with **Ignore Warnings** deselected in order to view any warnings displayed. You may then evaluate whether to ignore or rectify any warnings, clear the warning dialog, and then click **Save** again with the box selected to avoid stopping.
- ◆ **User Name:** User credential for login to the node.
- ◆ **Verify Mode:** Choose the verification mode as *Password* or *Key-based*, and then either enter your password or upload a private key file. If you choose *Key-based*, you must first enter a user name and then upload a private key file when connecting the node with a private key file.
- ◆ **Thinpool Device:** (conditional) Enter the Thinpool Device path that you configured for the master node (if applicable). For example: `/dev/mapper/docker-thinpool`. You must have already set up the Docker thin pool for all cluster nodes that need to use thinpools, as described in the *CDF Planning Guide*.
- ◆ **flannel IFace:** (conditional) Enter the flannel IFace value if the master node has more than one network adapter. This must be a single IPv4 address or name of the existing interface and will be used for Docker inter-host communication.

Click **Save**. Repeat the same for other master nodes.

- 12 Click **Next**.
- 13 (Conditional) For multi-node deployment, add additional worker nodes on the Add Worker Node page and click **Save**. To add a worker node click + (Add) and enter the required configuration information. Repeat this process for each of the worker nodes.
- 14 Click **Next**.
- 15 (Conditional) If you want to run the worker node on the master node, select **Allow suite workload to be deployed on the master node** and then click **Next**.

NOTE: Before selecting this option, ensure that the master node meets the system requirements specified for the worker node.

16 To configure each NFS volume, complete the following steps:

16a Navigate to the **File Storage** page.

16b For **File System Type**, select **Self-Hosted NFS**.

Self-hosted NFS refers to the **external NFS** that you created while preparing the environment for CDF installation.

16c For **File Server**, specify the IP address or FQDN of the NFS server.

16d For **Exported Path**, specify the following paths for the NFS volumes:

NFS Volume	File Path
arcsight-volume	<NFS_ROOT_FOLDER>/arcsight-vol
itom-vol-claim	<NFS_ROOT_FOLDER>/itom-vol
db-single-vol	<NFS_ROOT_FOLDER>/db-single-vol
itom-logging-vol	<NFS_ROOT_FOLDER>/itom-logging-vol
db-backup-vol	<NFS_ROOT_FOLDER>/db-backup-vol

16e Click **Validate**.

Ensure that you have validated all NFS volumes successfully before continuing with the next step.

17 Click **Next**.

18 To start deploying master and worker nodes, click **Yes** in the Confirmation dialog box.

19 Continue with uploading images to the local registry.

Uploading Images to the Local Registry

For the docker registry to deploy Fusion, it needs the following images associated with the deployment:

- ♦ `fusion-x.x.x.x`
- ♦ `analytics-x.x.x.x`

You must upload those images to the local registry as follows:

- 1** Launch a terminal session, then log in to the master node as `root` or a `sudo` user.
- 2** Change to the following directory:

```
cd /<cdf_installer_directory>/scripts/
```

For example:

```
cd /opt/fusion-installer-x.x.x.x/installers/cdf-x.x.x-x.x.x.x/scripts/
```

- 3** Upload the Analytics images to the local registry:


```
./uploadimages.sh -d <download_directory> -u registry-admin -p  
<cdf_password>
```

Example:

```
./uploadimages.sh -d /<download_directory>/fusion-installer-x.x.x.x/  
suite_images/analytics-x.x.x.x -u registry-admin -p <cdf_password>
```

4 Upload the Fusion image to the local registry:

```
./uploadimages.sh -d <download_directory> -u registry-admin -p  
<cdf_password>
```

Example:

```
./uploadimages.sh -d /<download_directory>/fusion-installer-x.x.x.x/  
suite_images/fusion-x.x.x.x -u registry-admin -p <cdf_password>
```

5 Continue with deploying Fusion.

Deploying Fusion

After you upload the images to the local directory, Container Deployment Foundation (CDF) uses these images to deploy the respective software in the cluster.

- 1 Switch back to the CDF Management Portal.
- 2 Click **Next** on the **Download Images** page because all the required packages are already downloaded and uncompressed.
- 3 After the **Check Image Availability** page displays All images are available in the registry, click **Next**.

If the page displays any missing image error, [upload the missing image](#).

- 4 After the Deployment of Infrastructure Nodes page displays the status of the node in green, click **Next**.

The deployment process can take up to 15 minutes to complete.

- 5 (Conditional) If any of the nodes show a red icon on the Deployment of Infrastructure Nodes page, click the **retry** icon.

IMPORTANT: CDF might display the red icon if the process times out for a node. Because the retry operation executes the script again on that node, ensure that you click **retry** only once.

- 6 After the Deployment of Infrastructure Services page indicates that all the services are deployed and the status indicates green, click **Next**.

The deployment process can take up to 15 minutes to complete.

(Optional) To monitor the progress of service deployment, complete the following steps:

6a Launch a terminal session.

6b Log in to the master node as `root`.

6c Execute the command:

```
watch 'kubectl get pods --all-namespaces'
```

- 7 Click **Next**.

- 8 Configure the pre-deployment settings in the CDF Management Portal, by making the following changes under **ANALYTICS**:
 - ◆ In the **Cluster Configuration** section, select **0** from the **Hercules Search Engine Replicas** drop-down list.

NOTE: By default, the value for **Hercules Search Engine Replicas** is **1**.

 - ◆ In the **Vertica Configuration** section, disable Vertica.
 - ◆ In the **Single Sign-on Configuration** section, specify the values for **Client ID** and **Client Secret**.
- 9 To finish the deployment, click **Next**.
- 10 Copy the Management portal link displayed on the **Configuration Complete** page.

Some of the pods on the **Configuration Complete** page might remain in a pending status until the product labels are applied on worker nodes.
- 11 (Conditional) For high availability and multi-master deployment, after the deployment has been completed, manually restart the keepalive process.
 - 11a Log in to the master node.
 - 11b Change to the following directory:

```
cd /<k8S_HOME>/bin/
```

For example:

```
cd /opt/arcsight/kubernetes/bin/
```
 - 11c Run the following script:

```
./start_lb.sh
```
- 12 Continue to the [post-installation steps](#).

9 Deploying ArcSight Fusion in an Existing Cluster

If you already have an existing cluster for Transformation Hub or ArcSight Investigate, you can deploy Fusion to the same cluster. Reusing existing clusters would reduce costs and system management effort compared to deploying this software in a new cluster.

- ◆ “Prerequisites” on page 51
- ◆ “Deploying Fusion to an Existing Cluster” on page 51


Prerequisites

Before installing Fusion, complete the following tasks:

- ◆ Review the [Installation Checklist](#) to understand the tasks involved for installing and configuring Fusion.
- ◆ Download the Fusion installation files and verify the signatures. You need the following files to deploy Fusion in an existing cluster:
 - ◆ `fusion-x.x.x.x`
 - ◆ `analytics-x.x.x.x` (if other products such as Investigate or Intersect are not deployed)
 - ◆ `arcsight-installer-metadata-x.x.x.x.tar`
- ◆ Upgrade the existing cluster to the correct version of the ArcSight Suite in the CDF Management Portal to deploy the current version of Fusion.

Deploying Fusion to an Existing Cluster

To deploy Fusion in an existing cluster, perform the following steps:

- 1 Log in to the CDF Management Portal.
- 2 Go to **Deployment** > **Deployments** >  > **Change**.
- 3 On the **Capabilities** page, select the following options:
 - ◆ Fusion
 - ◆ Analytics (if not selected already)
- 4 Click **Next** until you reach the **Import Container Images** page.
- 5 Launch a terminal session, then log in to the master node as `root` or as a `sudo` user.
- 6 Change to the following directory:

```
cd /<fusion_installer_directory>/scripts/
```

For example:

```
cd /opt/fusion-installer-x.x.x.x/installers/cdf-x.x.x-x.x.x.x/scripts/
```

- 7 (Conditional) If your deployment does not have Investigate/Identity Intelligence/Interset, upload the Analytics image to the local registry:

```
./uploadimages.sh -d /<download_directory>/fusion-installer-x.x.x.x/  
suite_images/analytics-x.x.x.x -u registry-admin -p <cdf_password>
```

- 8 Upload the Fusion image to the local registry:

```
./uploadimages.sh -d /<download_directory>/fusion-installer-x.x.x.x/  
suite_images/fusion-x.x.x.x -u registry-admin -p <cdf_password>
```

- 9 Switch to the CDF Management Portal, then click **CHECK AGAIN** to ensure that the images have been uploaded.
- 10 Click **Next** until you reach the **Configuration Complete** page.
- 11 After the **Configuration Complete** page displays all the pods in green, click **Next**.
- 12 Continue to the [post-installation steps](#).

10 Post-Installation Configuration

This chapter provides information about the post-installation configuration you must perform after deploying Fusion.

- ♦ [“Add Dashboard Roles and Permissions to the User Management Pod”](#) on page 53
- ♦ [“Labeling Nodes”](#) on page 53
- ♦ [“Connecting to an SMTP Server”](#) on page 55
- ♦ [“Integrating Fusion Single Sign-On with an External SAML 2 Identity Provider”](#) on page 55

Add Dashboard Roles and Permissions to the User Management Pod

Applies only when you deploy Fusion to an existing cluster

When you deploy Fusion to an existing cluster, the User Management pod does not automatically incorporate the new roles and permissions for the Dashboard to the **ADMIN** function in the user interface. Without the default roles, not even an administrative user can view the Dashboard. Deleting (restarting) the User Management pod adds the roles and permissions.

To delete the User Management pod:

- 1 Complete the Fusion deployment.
- 2 Enter the following command:

```
kubectl delete pod -n namespace management pod name
```

For example:

```
kubectl delete pod -n arcsight-installer-p2dlt hercules-management-7f876b4978-9xkl6
```

When you delete any pod, the pod will start automatically.

- 3 After the pod restarts, you can [log in to the Dashboard](#).

Labeling Nodes

IMPORTANT: *Does not apply if you used the `./install-single-node-post.sh` installation script*

Labeling a node tells Kubernetes what type of application can run on a specific node. It identifies application processing and qualifies the application as a candidate to run on a specific host system. Labeling is required only for worker nodes and not for master nodes.

The method for labeling nodes depends on whether you installed Fusion in a new or an existing cluster.

- ♦ [“Labeling Nodes in a New Cluster” on page 54](#)
- ♦ [“Labeling Nodes in an Existing Cluster” on page 54](#)

Labeling Nodes in a New Cluster

You can follow the instructions in this section to label the nodes manually for both single-node and multi-node deployment. However, for single node deployment, you can alternatively use the `/opt/<ArcSight_Fusion_Installer>/postinstall_label_master_node.sh` script (present in the scripts folder) to label the node automatically.

To label the nodes in a new cluster:

- 1 Open a new tab in a supported web browser.
- 2 Specify the URL for the CDF Management Portal:

```
https://Fusion_server:5443
```

Use the fully qualified domain name of the host that you specified in the **Connection** step during the CDF configuration. Usually, this is the master node’s FQDN.

- 3 Log in to the CDF Management Portal with the credentials of the administrative user that you provided during installation.
- 4 Select **Administration > Nodes**.
- 5 In **Predefined Labels**, click **+** to add labels.
- 6 Specify the following label:

```
analytics:yes
```

NOTE: Labels are case-sensitive. Ensure that you enter the values correctly.

- 7 (Conditional) For single-node deployment, drag all the newly-added labels to the worker node.
- 8 (Conditional) For multi-node deployment, drag and drop the new labels from the predefined set to each of the worker nodes based on your workload sharing configuration.

You might need to click **Refresh** to see the attached labels.

Labeling Nodes in an Existing Cluster

You do not need to label the nodes if you have Investigate because both Fusion and Investigate use the same label: `analytics:yes`. However, if your deployment does not have Investigate, you can either add Fusion to a new worker node or an existing worker node, if the node has adequate unused capacity.

- 1 Open a new tab in a supported web browser.
- 2 Specify the URL for the CDF Management Portal:

```
https://<Fusion-server>:5443
```

Use the fully qualified domain name of the host that you specified in the **Connection** step during the CDF configuration. Usually, this is the master node’s FQDN.

- 3 Log in to the CDF Management Portal with the credentials of the administrative user that you provided during installation.
- 4 Select **Administration > Nodes**.
- 5 (Conditional) To deploy Fusion on a new node, click **ADD**.
- 6 Specify the details of the node, such as worker node type, host name or IP address, user name, and password.
- 7 Drag the `analytics:yes` label to the newly added worker node or an existing worker node.

Connecting to an SMTP Server

To ensure that Fusion users receive email notifications, configure the connection to your SMTP server. For example, if you do not use SAML 2 authentication, users will need notifications to help reset their forgotten passwords.

- 1 Log in to the CDF Management Portal with the credentials of the administrative user that you provided during installation.
- 2 Select **ANALYTICS**.
- 3 Under **User Management Configuration**, configure the SMTP settings.

Integrating Fusion Single Sign-On with an External SAML 2 Identity Provider

You can configure Fusion to establish a trust relationship with your external identity provider. With this authentication method, a user's email address, specified in the 'email' claim value from the SAML2 Identity Provider, maps to the userID for the Fusion user.

NOTE

- ♦ You should time-synchronize Fusion and the external SAML 2.0 IDP to the same NTP server. In the configuration UI, ensure that the session timeout matches the same value that the external IDP has configured for user session timeouts.
- ♦ **Regarding the Trusted Provider Metadata**, the metadata document for a trusted SAML provider with which an SSO-defined provider interacts must be obtained in a provider-specific manner. While not all providers do so, many supply their metadata documents via URL.

Once the trusted provider's metadata document (or the URL-accessible location of the document) is obtained, you must configure the SSO provider that will interact with the trusted provider with the trusted provider's metadata. In the document, modify the `<Metadata>` element within the `<AccessSettings>` element under either the `<TrustedIDP>` element or the `<TrustedSP>` element. For example:

```
com.microfocus.sso.default.login.saml2.mapping-attr = email
```

The `email` attribute refers to the email attribute name from the SAML2 IDP.

To integrate with an external SAML provider:

1 On the CDF server, open the `sso-configuration.properties` file, located by default in the `/opt/arcsight/nfs/vol/arcsight/sso/default` directory.

2 Add the following properties to the file:

- ◆ `com.microfocus.sso.default.login.method = saml2`
- ◆ `com.microfocus.sso.default.saml2.enabled = true`

3 To specify the address where the IDP supplies its metadata document, complete one of the following actions:

- ◆ Add the following property to the file:

```
com.microfocus.sso.default.login.saml2.metadata-url = IDP SAML
metadata URL
```

For example, a Keycloak server URL could be `https://KeycloakServer/auth/realms/YourRealm/protocol/saml/descriptor`.

NOTE: For HTTPS to work properly, you must import the IDP certificates to the Fusion single sign-on keystore as described in [Step 5](#).

- ◆ Convert the metadata xml file to a base64 string, then add the following property to the file:

```
com.microfocus.sso.default.login.saml2.metadata = base64 encoded
metadata xml
```

4 Save the changes to the `sso-configuration.properties` properties file.

5 (Conditional) If you specified the metadata URL in [Step 3](#), complete the following steps to import the IDP certificate to the SSO keystore:

5a Copy the IDP certificate to following location:

```
/path/to/sso/default/
```

5b Get the pod information using the following command:

```
kubectl get pods --all-namespaces | grep osp
```

5c Open a terminal in the currently running `hercules-osp` pod:

```
kubectl exec -it hercules-osp-xxxxxxxxxx-xxxxx -n arcsight-
installer-xxxxx -c hercules-osp -- bash
```

5d Import the IDP certificate file using the following commands:

5d1 `cd /usr/local/tomcat/conf/default/`

5d2 `keytool -importcert -file CertificateFileName -keystore sso.keystore -storepass $KEYSTORE_PASSWORD -alias AliasName`

where

- ◆ **CertificateFileName** represents the name of the certificate file that you want to import.
- ◆ **AliasName** represents the new alias name that you want to assign to the certificate in the SSO keystore.

6 Restart the pod by completing the following steps:

6a To get the pod information, enter the following command:


```
kubectl get pods --all-namespaces | grep osp
```

6b To delete the currently running pod, enter the following command:

```
kubectl delete pod hercules-osp-xxxxxxxxxx-xxxxx -n arcsight-  
installerxxxxx
```

7 Retrieve the Fusion SSO SAML service provider metadata from the Fusion server:

```
https://Fusion_server/osp/a/default/auth/saml2/spmetadata
```

where **Fusion_server** represents the host name of the Fusion server.

8 Use the Fusion SSO SAML service provider metadata to configure your IDP.

For more information, see the IDP software documentation.

9 To establish a trust relationship between Fusion SSO and your IDP software, create certificates for your IDP software.

For more information on how to create and import certificates in your IDP software, see the IDP software documentation.

11 Verifying the Installation

To determine whether the installation is successful, launch the Dashboard. The first time that you connect after deployment, the application prompts you to create credentials for the product administrator.

- 1 Open a supported web browser.
- 2 Specify the URL for the Dashboard:

```
https://<Fusion-server>/dashboard
```

- 3 Specify credentials for the appropriate administrator, using the following format:
name@domain.com
- 4 Specify credentials for the appropriate administrator.

NOTE: If your environment includes Intersect, review the following requirements for the initial administrative user:

- ◆ For the first-time login, you must enter credentials for the *Intersect administrator*, which is the administrative account that was specified during the deployment of Intersect. Otherwise, Intersect users will not have single sign-on access to the Dashboard.
- ◆ The Intersect administrator account cannot access analytics data in Intersect because the application automatically assigns it to the *Administrative Tenant* role. So, this role cannot view content in widgets that display data from Intersect.
- ◆ Other users created with Dashboard permissions automatically receive the *Default Tenant* status in Intersect. These users can view content in widgets that display data from Intersect.

-
- 5 Log in with the new credentials.
 - 6 (Conditional) To connect with ESM, continue to [Chapter 12, “Integrating Data and Users from Enterprise Security Manager,”](#) on page 63.
 - 7 (Conditional) To connect with Intersect, continue to [Chapter 13, “Integrating Data from Intersect,”](#) on page 67.
 - 8 Create additional [Dashboard users](#).



Configuring Fusion

This section provides guidance for configuring Fusion data sources and users after a successful installation.

- ♦ [Chapter 12, “Integrating Data and Users from Enterprise Security Manager,” on page 63](#)
- ♦ [Chapter 13, “Integrating Data from Interset,” on page 67](#)
- ♦ [Chapter 14, “Adding Users and Groups to the Dashboard,” on page 69](#)

12 Integrating Data and Users from Enterprise Security Manager

Fusion allows you to integrate data from Enterprise Security Manager (ESM). With single sign-on enabled, ESM and Fusion users can easily access the ArcSight Console, ArcSight Command Center, Fusion, and the [REST APIs](#). As a convenience, you can import users that are already authorized for ESM.

- ♦ [“Understanding How ESM Users Access Fusion” on page 63](#)
- ♦ [“Importing Users from ESM” on page 63](#)
- ♦ [“Enabling Single Sign-on with ESM” on page 64](#)
- ♦ [“Integrating Data from Enterprise Security Manager” on page 65](#)

Understanding How ESM Users Access Fusion

Rather than manually adding users in Fusion, we recommend that you create the users in ESM then [import](#) them. For the imported ESM users to log in to Fusion and be able to access ESM data, the following conditions apply:

- ♦ You must [enable single sign-on \(SSO\) access](#) for ESM and Fusion users.
- ♦ Users must have an account in both ESM and Fusion.
- ♦ You must configure the **External User ID** and **E-mail** fields in the ESM accounts to comply with the *name@domain.com* format.
- ♦ Users must log in to Fusion with the **External User ID** from their ESM account.
- ♦ If your environment does not use SAML 2 authentication, ensure that you have [configured the SMTP server settings](#) for Fusion. Users might need to set a password for Fusion the first time that they log in. This requires those users to initiate the “Forgot Password” function and receive an email notification.

Importing Users from ESM

You can import users that are already authorized for ESM. You need to have at least one [role](#) available in the Dashboard to assign to these users.

- 1 In the ArcSight Console, ensure that the **External User ID** and **E-mail** fields for each account comply with the following format:

name@domain.com

- 2 Log in to Fusion.
- 3 Click **ADMIN > Account Groups > Import Users**.

- 4 Select the role that you want to assign to the imported users.
- 5 Select **IMPORT USERS**.

As you add more users to ESM, you can run the import process again. Fusion ignores duplicates of user accounts that have been imported previously.

Enabling Single Sign-on with ESM

You must configure ESM to use **OSP Client Only Authentication**. If your ESM environment currently uses an external SAML 2 client authentication, you must [delegate the Fusion SSO provider](#) to connect to the SAML client. If you do not use SAML 2 authentication, you will need to [configure SMTP settings](#) for Fusion to support forgotten password activity.

This procedure assumes that you have already installed or upgraded ESM.

- 1 Change the authentication settings for the ESM Manager service:
 - 1a On the ESM server, start the configuration wizard by entering the following command from the `/opt/arcsight/manager/bin/` directory:

```
arcsight managersetup -i console
```
 - 1b Advance through the wizard until you reach the authentication settings.
 - 1c Select **OSP Client Only Authentication**, then click **Next**.
 - 1d To specify the host and port for the OSP server, use the following format:

```
domain_name:port
```
 - 1e To specify the host and port for the ArcSight Command Center, use the following format:

```
domain_name:port
```

Typically, the host and port are the same as those for the ArcSight Manager.
 - 1f Specify a **Tenant Name for OSP**. The default value is `default`.
 - 1g Click **Next** until you complete your changes in the wizard.
 - 1h Restart the ESM Manager service using the following commands:

```
/etc/init.d/arcsight_services stop manager  
/etc/init.d/arcsight_services start manager
```
- 2 Change the authentication settings for the ArcSight Console (Console):
 - 2a From the Console's `/bin` directory, enter one of the following commands:
On Windows: `arcsight.bat consolesetup`
On Linux: `./arcsight consolesetup`
 - 2b Advance through the wizard until you reach the authentication settings.
 - 2c Select **OSP Client Only Authentication**.
 - 2d Click **Next** until you complete your changes in the wizard.
- 3 Configure the SSO settings in the CDF Management Portal:
 - 3a Connect to the Portal:

```
https://Fusion_server:5443
```


- 3b** Log in with the credentials of the administrative user that you provided during installation.
- 3c** Select **ANALYTICS**.
- 3d** Under **Single Sign-on Configuration**, specify the **Client ID** and **Client Secret**.
- 3e** Under **ArcSight ESM Host Configuration**, verify the settings for the ESM host and port that were specified during deployment.
- 4** (Conditional) To use an external SAML2 authentication method, continue to [“Integrating Fusion Single Sign-On with an External SAML 2 Identity Provider” on page 55](#).
- 5** (Conditional) If you do not use an external SAML 2 authentication method, ensure that users can receive email notifications to change their Fusion password. Continue to [“Connecting to an SMTP Server” on page 55](#).


Integrating Data from Enterprise Security Manager

To view ESM data in the Dashboard, update the settings in the CDF Management Portal.

- 1 Open a new tab in a supported web browser.
- 2 Specify the URL for the CDF Management Portal:

```
https://Fusion_server:5443
```

Use the fully qualified domain name of the host that you specified in the **Connection** step during the CDF configuration. Usually, this is the master node’s FQDN.

- 3 Login to the **CDF Management Portal** with the credentials of the administrative user that you provided during installation.
- 4 In the **CDF Management Portal**, click the  icon for the installed suite.
- 5 Select **Reconfigure**.
- 6 On the **Configuration** page, select **ANALYTICS**.
- 7 In the **ArcSight ESM Host Configuration** section, complete the following steps:
 - 7a** For **ESM host**, specify the fully-qualified host name or IP address of the server that hosts ESM.
 - 7b** For **ESM port**, specify the port associated with the ESM host. The default value is 8443.

13 Integrating Data from Intersect

Fusion supports Intersect within the same cluster. As an administrator, you do not have to perform any additional configurations in Fusion to integrate data from Intersect.

Intersect users automatically have single sign-on access as long as the [first administrative account](#) for Fusion matches the credentials for the *Intersect administrator*, which is the administrative account that was specified during the deployment of Intersect. Users with [Dashboard permissions](#) automatically receive the *Default Tenant* status in Intersect. These users can view content in widgets that display data from Intersect. The [User Guide for ArcSight Fusion](#), available as context-sensitive help within Fusion, indicates whether a widget needs data from a specific data source.

14 Adding Users and Groups to the Dashboard

The Dashboard allows you to incorporate users either by manually adding them or by [importing users and groups from ESM](#). To assign permissions to these users, you can create roles with specific sets of permissions and add users to those roles.

For more information about assigning permissions to users and roles, see the *User Guide for ArcSight Fusion* embedded in the product or posted with the [documentation for ArcSight Fusion](#).

IV

Managing ArcSight Fusion

This section provides information about managing ArcSight Fusion.

- ♦ [Chapter 15, “Configuring the Dashboard,” on page 73](#)
- ♦ [Chapter 16, “Using REST APIs with Fusion,” on page 75](#)
- ♦ [Chapter 17, “Restarting Nodes in the Cluster,” on page 77](#)
- ♦ [Chapter 18, “Resetting the CDF Administrator Password,” on page 79](#)
- ♦ [Chapter 19, “Renewing CDF Certificates,” on page 81](#)
- ♦ [Chapter 20, “Creating and Adding CDF Certificate Authority,” on page 83](#)

15 Configuring the Dashboard

Fusion allows you to perform the following configuration:


- ◆ Specify the log level for the Dashboard Metadata Web Application service.
- ◆ Specify the maximum number of days of data that the Dashboard can collect and display. The default value is **365 days**.

To configure your dashboard:

- 1 Open a new tab in a supported web browser.
- 2 Specify the URL for the CDF Management Portal:

```
https://<Fusion-server>:5443
```

Use the fully-qualified domain name of the host that you specified in the **Connection** step during the CDF configuration. Usually, this is the master node's FQDN.

- 3 Log in to the **CDF Management Portal** with the credentials of the administrative user that you provided during installation.
- 4 In the CDF Management Portal, click  and select **Reconfigure**.
- 5 On the **Configuration** page, click the **ArcSight Fusion** tab.
- 6 On the **ArcSight Fusion** tab, under the Dashboard Configuration section, select the log level from the **Metadata Log Level** drop-down list. The default value is **info**.
- 7 In **Maximum Time Range (Days)**, specify the maximum number of days of data that the dashboard can collect and display. The default value is 365 days.

16 Using REST APIs with Fusion

The Fusion dashboard user interface uses REST APIs to manage dashboard configurations and these APIs can be accessed directly. For example, you might want to update a particular user's dashboard. You can access the *Dashboard Metadata REST API Documentation* for the Dashboard at the following URL:

```
https://<Fusion Server>/metadata/rest-api-docs
```

NOTE: For [single sign-on access](#) to the REST APIs, specify the values for **Client ID** and **Client Secret**, in the **Single Sign-on Configuration** section. For more information, see [step 9 in the Deploying Fusion section](#).

17 Restarting Nodes in the Cluster

If you want to restart or shut down any node in the Fusion cluster, you must stop the Kubernetes services running on the node. Otherwise, the Kubernetes pods will not start after the restart.

You can restart nodes in one of the following ways:

- ♦ [“Restarting Nodes by Using Scripts” on page 77](#)
- ♦ [“Restarting Nodes Manually” on page 77](#)

Restarting Nodes by Using Scripts

IMPORTANT: This information is applicable only if you have installed Fusion by using scripts.

To restart the node:

- 1 Log in to the node you need to restart.
- 2 To restart the node execute the following command:

```
/opt/fusion/bin/single-node-util.sh reboot_node
```

Restarting Nodes Manually

IMPORTANT: This information is applicable only if you have installed Fusion manually.

To restart the node manually:

- 1 (Conditional) If the node contains CDF, perform the following steps:

- 1a Log in to the node you need to restart as `root`.
- 1b Change to the following directory:

```
cd <K8S_HOME>/bin/
```

For example:

```
/opt/arcsight/kubernetes/bin
```

- 1c Stop the kubernetes services by using the following command:

```
kube-stop.sh
```

- 1d Unmount Kubernetes volumes by using the following command:

```
kubelet-umount-action.sh
```

- 2 Restart the node using the following command:

```
reboot
```

3 (Conditional) After the node restarts, perform the following steps if the node contains CDF:

3a Log in to the node as `root`.

3b Change to the following directory:

```
cd <K8S_HOME>/bin/
```

For example:

```
/opt/arcsight/kubernetes/bin
```

3c Check whether all Kubernetes services are running:

```
kube-status.sh
```

3d (Conditional) If any service is not running, start the service by using the command:

```
kube-start.sh
```

18 Resetting the CDF Administrator Password

To change the administrator password of the CDF Management Portal, perform the following steps:

- 1 Log in to the CDF Management Portal (<https://<Fusion Server>:5443>) using the administrator user ID and password that you specified during the CDF installation.
- 2 On the main page, click **Application** > **IdM Administration**. And then, click the **SRG** image.
- 3 In **Organization**, click **Users**.
- 4 In the list of users on the right, select **Admin** and click **Edit**.
- 5 Click **Remove Password**.
- 6 Click **Add Password**.
- 7 Enter a new admin password, then click **Save**.

19 Renewing CDF Certificates

The validity of a CDF certificate is one year. If you do not upgrade the CDF certificate within one year, the certificate will expire and you must renew the certificate. The CDF certificates can be renewed before and after expiration.

CDF contains the following certificates:

- ♦ Internal certificates, which are used within the cluster nodes, such as `client.crt`, `client.key`, `server.crt`, `server.key`, `kubernetes.crt`, and `kubernetes.key`.
- ♦ External certificates, which are used for the ingress service of the CDF Management Portal.

Renewing Certificates Before Expiration

You can renew both internal and external certificates before expiration.

To renew certificates before expiration:

1 Log in to the master node.

2 Change to the following directory:

```
cd <k8s_HOME>
```

By default, `k8s_HOME` is `/opt/kubernetes`.

3 (Conditional) For internal certificates, run the following command to generate new certificates:

```
./scripts/renewCert --renew -t internal
```

In a multi-node deployment, executing the above command automatically distributes the new certificates to all nodes in the cluster.

4 (Conditional) For external certificates, run the following command to generate new certificates:

```
./scripts/renewCert --renew -t external
```

Renewing Certificates After Expiration

You can renew both internal and external certificates after expiration.

To renew certificates after expiration:

1 Log in to the master node.

2 Change to the following directory:

```
cd <k8s_HOME>
```

By default, `k8s_HOME` is `/opt/kubernetes`.

3 (Conditional) For internal certificate:

3a Run the following command to generate new `client.crt`, `client.key` and `server.crt` certificates:

```
./scripts/renewCert --renew -V 365 -t internal
```

3b (Conditional) If you have multiple master nodes, run the following on all the master nodes:

```
./scripts/renewCert --renew -t internal
```

4 (Conditional) For external certificates, run the following command:

- ◆ To generate new external self-signed certificates:

```
./scripts/renewCert --renew -t external
```

- ◆ To generate the external custom self-signed certificates:

```
./scripts/renewCert --renew -t external --tls-cert /<cert file  
directory>/<cert file> --tls-key <private key directory>/<private  
key> [--tls-cacert <CA cert directory>/<CA cert file>]
```

20 Creating and Adding CDF Certificate Authority

The cluster maintains its own certificate authority (CA) to issue certificates for external communication. A self-signed CA is generated during the installation of CDF by default. Pods of the deployed products use the certificates generated by the CA on pod startup. When configuring SSL, you must create a new CA and add the CA to CDF, which will be used by all the products in the cluster.

IMPORTANT: If you change the CA after Fusion deployment, you will have to uninstall and reinstall the CDF suite. Uninstalling the CDF suite will uninstall all the installed capabilities such as Fusion and Analytics. We recommend that you perform this procedure when Fusion is first installed to avoid downtime and data loss.

To create a new CA and add the CA to CDF, perform the following:

1 Create a new CA by performing the following steps:

1a Create a directory and configure the directory permissions using the following command:

```
mkdir /root/ca
cd /root/ca
mkdir certs crl newcerts private
chmod 700 private
touch index.txt
echo 1000 > serial
```

1b Open the configuration file in a text editor (`vi /root/ca/openssl.cnf`) and add the following content (values shown here are examples; change parameter values to match your's):

```
# OpenSSL root CA configuration file.
# Copy to `/root/ca/openssl.cnf`.
[ ca ]
default_ca = CA_default
[ CA_default ]
# Directory and file locations.
dir = /root/ca
certs = $dir/certs
crl_dir = $dir/crl
new_certs_dir = $dir/newcerts
database = $dir/index.txt
serial = $dir/serial
RANDFILE = $dir/private/.rand
# The root key and root certificate.
private_key = $dir/private/ca.key.pem
certificate = $dir/certs/ca.cert.pem
# For certificate revocation lists.
```

```

crlnumber = $dir/crlnumber
crl = $dir/crl/ca.crl.pem
crl_extensions = crl_ext
default_crl_days = 30
# SHA-1 is deprecated, so use SHA-2 instead.
default_md = sha256
name_opt = ca_default
cert_opt = ca_default
default_days = 375
preserve = no
policy = policy_strict
[ policy_strict ]
# The root CA should only sign intermediate certificates that match.
# See the POLICY FORMAT section of `man ca`.
countryName = match
stateOrProvinceName = match
organizationName = match
organizationalUnitName = optional
commonName = supplied
emailAddress = optional
[ policy_loose ]
# Allow the intermediate CA to sign a more diverse range of
certificates.
# See the POLICY FORMAT section of the `ca` man page.
countryName = optional
stateOrProvinceName = optional
localityName = optional
organizationName = optional
organizationalUnitName = optional
commonName = supplied
emailAddress = optional
[ req ]
# Options for the `req` tool (`man req`).
default_bits = 2048
distinguished_name = req_distinguished_name
string_mask = utf8only
# SHA-1 is deprecated, so use SHA-2 instead.
default_md = sha256
# Extension to add when the -x509 option is used.
x509_extensions = v3_ca
[ req_distinguished_name ]
countryName = Country
stateOrProvinceName = State
localityName = Locality
0.organizationName = EntCorp
organizationalUnitName = OrgName
commonName = Common Name
emailAddress = Email Address
# Optionally, specify some defaults.
countryName_default = <your country code>
stateOrProvinceName_default = <your state or province>
localityName_default =
0.organizationName_default = <your company name>
organizationalUnitName_default =
emailAddress_default =

```

```

[ v3_ca ]
# Extensions for a typical CA (`man x509v3_config`).
subjectKeyIdentifier = hash
authorityKeyIdentifier = keyid:always,issuer
basicConstraints = critical, CA:true
keyUsage = critical, digitalSignature, cRLSign, keyCertSign
[ v3_intermediate_ca ]
# Extensions for a typical intermediate CA (`man x509v3_config`).
subjectKeyIdentifier = hash
authorityKeyIdentifier = keyid:always,issuer
basicConstraints = critical, CA:true, pathlen:0
keyUsage = critical, digitalSignature, cRLSign, keyCertSign
[ usr_cert ]
# Extensions for client certificates (`man x509v3_config`).
basicConstraints = CA:FALSE
nsCertType = client, email
nsComment = "OpenSSL Generated Client Certificate"
subjectKeyIdentifier = hash
authorityKeyIdentifier = keyid,issuer
keyUsage = critical, nonRepudiation, digitalSignature,
keyEncipherment
extendedKeyUsage = clientAuth, emailProtection
[ server_cert ]
# Extensions for server certificates (`man x509v3_config`).
basicConstraints = CA:FALSE
nsCertType = server
nsComment = "OpenSSL Generated Server Certificate"
subjectKeyIdentifier = hash
authorityKeyIdentifier = keyid,issuer:always
keyUsage = critical, digitalSignature, keyEncipherment
extendedKeyUsage = serverAuth
[ crl_ext ]
# Extension for CRLs (`man x509v3_config`).
authorityKeyIdentifier=keyid:always
[ ocspr ]
# Extension for OCSP signing certificates (`man ocspr`).
basicConstraints = CA:FALSE
subjectKeyIdentifier = hash
authorityKeyIdentifier = keyid,issuer
keyUsage = critical, digitalSignature
extendedKeyUsage = critical, OCSPSigning

```

1c Generate a CA root key:

```

cd /root/ca
openssl genrsa -out private/ca.key.pem 4096
chmod 400 private/ca.key.pem

```

1d Create a CA cert:

```

openssl req -config openssl.cnf -key private/ca.key.pem -new -x509 -
days 375 -sha256 -extensions v3_ca -out certs/ca.cert.pem

```

1e Verify the root CA:

```

chmod 444 certs/ca.cert.pem
openssl x509 -noout -text -in certs/ca.cert.pem

```

2 Add the CA by running the following command:

```
<K8S_HOME>/scripts/cdf-updateRE.sh write --re-key=private/ca.key.pem --  
re-crt=certs/ca.cert.pem --re-ca=certs/ca.cert.pem
```

3 To verify, read the CDF CA file by executing the following command and ensure that it is same as the `ca.cert.pem` file. You must execute the following command on the initial master node:

```
<K8S_HOME>/scripts/cdf-updateRE.sh read
```