

Micro Focus ArcSight Platform Technical Requirements Guide

Technical Requirements for the ArcSight Platform

This document describes the system, software, and OS technical guidelines for effective operations and administration of the ArcSight Platform 23.2. The Platform enables you to deploy a combination of security, user, and entity solutions into a single cluster within the OPTIC Management Toolkit (OMT) environment. The core services for this OMT environment, including the Dashboard and user management, are provided by a common layer called Fusion.

- ["Software Requirements" on page 4](#)
- ["File System Options" on page 7](#)
- ["Data Types Associated with Connectors Supported by Intelligence" on page 9](#)
- ["Firewall Ports" on page 11](#)
- ["Hybrid Cloud Support" on page 11](#)
- ["Examples of Deployment Scenarios" on page 13](#)
- ["System Performance Benchmarks for Sizing and Tuning" on page 23](#)



To determine the best path for upgrading your environment to this release, see [ArcSight Upgrade Paths](#) on the ArcSight documentation site.



Customers running on platforms not listed in this document or with untested configurations will be supported until the point Micro Focus determines the root cause is the untested platform or configuration. According to the standard defect-handling policies, Micro Focus will prioritize and fix issues we can reproduce on the tested platforms.

Additional Documentation

The ArcSight Platform documentation library includes the following resources.

- [Release Notes for ArcSight Platform 23.2](#), which provides an overview of the products deployed in this suite and their latest features or updates.
- [Administrator's Guide for ArcSight Platform 23.2](#), which contains installation, user, and deployment guidance for the ArcSight software products and components that you deploy in the containerized platform.
- [User's Guide for Fusion in the ArcSight Platform](#), which is embedded in the product to provide both context-sensitive Help and conceptual information.

- [ArcSight Platform Upgrade Paths](#), which provides guidance for determining the upgrade path for your ArcSight Platform version.
- [Product Support Lifecycle Policy](#), which provides information on product support policies.

Software Requirements

This section lists the software needed to install and run the ArcSight Platform.

- ["Operating Systems" below](#)
- ["Software Components" on page 6](#)

For operating systems and SmartConnectors, the **Certified** label indicates our preferred platform; MicroFocus has performed extensive testing and has confidence that it will run without issues. The **Supported** label indicates an acceptable platform; MicroFocus has tested the software and has confidence that it will run without issues.

Operating Systems

Your environment must have at least a minimal installation of the specified operating system.

Category	Requirement
Database cluster nodes	<p><i>A FIPS-compliant version of the operating systems:</i></p> <p>Certified - Red Hat Enterprise Linux 8.6 (x86, x64)</p> <p>Supported - Red Hat Enterprise Linux 8.8 (x86, x64)</p>
Master/Worker cluster nodes	<p>Certified</p> <ul style="list-style-type: none"> • Red Hat Enterprise Linux 8.6 (x86, x64) • Red Hat Enterprise Linux 7.9 (x86, x64) • CentOS 7.9 (x86, x64) <p>Supported - Rocky Linux 8.6 (x86, x64)</p>
ESM Server	<p>Certified</p> <ul style="list-style-type: none"> • Red Hat Enterprise Linux 8.6 (x86, x64) • Red Hat Enterprise Linux 7.9 (x86, x64) • CentOS 7.9 (x86, x64) • SUSE Linux Enterprise Server 15 Service Pack 3 <p>Supported - Rocky Linux 8.6 (x86, x64)</p>

Category	Requirement
ESM Console	<p>Certified</p> <ul style="list-style-type: none"> • Red Hat Enterprise Linux 8.6 (x86, x64) • Red Hat Enterprise Linux 7.9 (x86, x64) • CentOS 7.9 (x86, x64) • SUSE Linux Enterprise Desktop 15 Service Pack 4 • macOS Big Sur, Monterrey, Ventura • Windows Server 2019 • Windows 10 Enterprise (including patches) <p>Supported - Rocky Linux 8.6 (x86, x64)</p>
Active-Passive High Availability (APHA)	<p>Certified</p> <ul style="list-style-type: none"> • Red Hat Enterprise Linux 8.6 (x86, x64) • Red Hat Enterprise Linux 7.9 (x86, x64) • CentOS 7.9 (x86, x64) • SUSE Linux Enterprise Server 15 Service Pack 3 <p>Supported - Rocky Linux 8.6 (x86, x64)</p>
Connectors	<p>Certified</p> <ul style="list-style-type: none"> • Red Hat Enterprise Linux 8.6 (x86, x64) • Red Hat Enterprise Linux 7.9 (x86, x64) • CentOS 7.9 (x86, x64) • SUSE Linux Enterprise Server 15 Service Pack 4 • Windows Server 2022 <p>Supported</p> <ul style="list-style-type: none"> • Rocky Linux 8.6 (x86, x64)
ArcMC (non-containerized)	<p>Certified</p> <ul style="list-style-type: none"> • Red Hat Enterprise Linux 8.6 (x86, x64) • Red Hat Enterprise Linux 7.9 (x86, x64) • CentOS 7.9 (x86, x64) <p>Supported - Rocky Linux 8.6 (x86, x64)</p>

Software Components

Category	Requirement
File systems	One of the following: <ul style="list-style-type: none">• EXT4 (recommended)• EXT3• Logical Volume Manager (LVM)• XFS
Data Collection	Certified - SmartConnector 8.4 Supported - SmartConnector 8.2 or later
Browser	<ul style="list-style-type: none">• Google Chrome• Microsoft Edge (Chromium-based)• Mozilla Firefox <p>Browsers should not use a proxy to access the OPTIC Management Toolkit (OMT) applications in the containerized ArcSight Platform because this might result in inaccessible web pages.</p>

File System Options

This section describes the available file system options.

- ["Network File System Options" below](#)
- ["Network File System \(NFS\) Minimum Directory Sizes" below](#)
- ["Install File System Options" on the next page](#)

Network File System Options

The following table lists the minimum network file system (NFS) options.

Category	Minimum Requirement
NFS Types	<ul style="list-style-type: none"> • Amazon EFS • HPE 3PAR File Persona • Linux-based NFS • NetApp
NFS Server Versions	<ul style="list-style-type: none"> • NFSv4.1 <p>(You might enable additional versions on the NFS server, however ArcSight Platform only uses NFSv4.1.</p>

Network File System (NFS) Minimum Directory Sizes

The following table lists the minimum required size for each of the NFS installation directories.

Directory	Minimum Size	Description
{NFS_ROOT_DIRECTORY}/itom-vol	130 GB	This is the OMT NFS root folder, which contains the OMT database and files. The disk usage will grow gradually.
{NFS_ROOT_DIRECTORY}/db-single-vol	Start with 10 GB	This volume is only available when you did not choose PostgreSQL High Availability (HA) for OMT database setting. It is for OMT database. During the install you will not choose the Postgres database HA option.

{NFS_ROOT_DIRECTORY}/db-backup-vol	Start with 10 GB	This volume is used for backup and restore of the OMT Postgres database. Its sizing is dependent on the implementation's processing requirements and data volumes.
{NFS_ROOT_DIRECTORY}/itom-logging-vol	Start with 40 GB	This volume stores the log output files of OMT components. The required size depends on how long the log will be kept.
{NFS_ROOT_DIRECTORY}/arcsight-volume	10 GB	This volume stores the component installation packages.

Install File System Options

The following table lists the supported file system options.

Category	Requirement
ext4	—
xfs	ftype=1

Data Types Associated with Connectors Supported by Intelligence

This section describes the data types associated with the SmartConnectors and FlexConnector types that Intelligence supports.

Data Type	Supported Smart Connectors
Access	SmartConnector for Microsoft Windows Event Log – Native Application and System Event Support SmartConnector for Microsoft Windows Event Log – Unified Application and System Event Support
Active Directory	SmartConnector for Microsoft Windows Event Log – Native Application and System Event Support
VPN	SmartConnector for Microsoft Network Policy Server File SmartConnector for Pulse Secure Pulse Connect Secure Syslog SmartConnector for Citrix NetScaler Syslog SmartConnector for Nortel Contivity Switch Syslog
Web Proxy	SmartConnector for Microsoft Forefront Threat Management Gateway File SmartConnector for Squid Web Proxy Server File SmartConnector for Blue Coat Proxy SG Multiple Server File
Repository	FlexConnector Type - ArcSight FlexConnector Regex File

Additional Considerations

Consider the following:

- A fuller set of SmartConnectors is supported for those sources that provide relevant data to the Intelligence analytics models. Micro Focus might need to examine sample logs to optimize analysis of data from this broader set of sources.
- For supported data types, Intelligence provides support for new devices that provide relevant data to the Intelligence analytics models. For more information, see "[Adding Support for New Devices](#)" in the *Administrators Guide for ArcSight Platform*.
- Intelligence supports the SmartConnectors listed. However, additional capabilities you might deploy, such as Recon, might support a wider set of SmartConnectors/FlexConnector types.

- Micro Focus advises against configuring event aggregation for data to be processed by ArcSight Intelligence. If you wish to use ArcSight Intelligence with aggregated events, contact [Micro Focus Customer Support](#).

Firewall Ports

When you install ArcSight Platform and deploy the associated capabilities, you will need to open firewall ports for many of the elements that make up the Platform. For more information about the ports for each component, see "[Understanding Firewall Ports for the ArcSight Platform](#)" in the *Administrator's Guide to ArcSight Platform 23.2*.


Hybrid Cloud Support

The ArcSight Platform can be deployed to a variety of locations, including on-premises and cloud. The [OMT infrastructure](#), the [capabilities](#) deployed on it, and the [ArcSight Database](#) have been tested when deployed together to the same location. For these components (which are unshaded in Figure 1), the following table specifies the combination of location, [Kubernetes service](#) and [ArcSight Database Communal Storage](#) service that has been tested, whereas other combinations have not been tested.

For the Database, you can choose one of several S3-compatible object store technologies. Note that not all S3-compatible object store technologies provide the same performance or capabilities. Research is important to determine the ideal solution for your needs. We've found that MinIO works well in the testing that we performed in our labs, but other solutions might work better for your environment. To help you get an idea of how to configure MinIO for use with ArcSight Database, see the example in [Configuring the Database for MinIO Storage](#) in the *Administrator's Guide to ArcSight Platform*.

[Other related components](#), such as an SMTP server, can be deployed in a hybrid cloud manner, with some deployed to different locations. To understand the tested scenarios for these components, see the documentation related to that component.

Figure 1. Unshaded components indicate those affected by the setting in the table



! Customers running on platforms not specified in this document or with untested configurations will be supported until the point Micro Focus determines that the root cause is the untested platform or configuration. According to the standard defect-handling policies, Micro Focus will prioritize and fix issues we can reproduce on the tested platforms.

Deployment Location	Kubernetes	Communal Storage
Amazon Web Services (AWS)	Elastic Kubernetes Service (EKS) <ul style="list-style-type: none"> v1.24* v1.25 v1.26 	AWS S3
Azure	Azure Kubernetes Service (AKS) <ul style="list-style-type: none"> v1.24 v1.25* v1.26 	Azure Blob
On-premises	OMT embedded	"bring-your-own" On-premises S3-compatible storage

*Verified in our test environments. However, we will support the other versions listed here.

Examples of Deployment Scenarios

You can deploy the ArcSight Platform capabilities in a variety of ways. The most basic deployment option is an all-in-one system that contains a limited number of capabilities on a single node. The single-node deployment is suitable for small workloads or to use as a proof-of-concept environment. For large workloads, you will need a multi-node environment, possibly with multiple masters. There are many scenarios and considerations involved in creating your environment. Please see "[Prerequisites and Considerations for Adding Capabilities](#)" in the *Administrator's Guide for ArcSight Platform*.

This section provides some examples on how you could deploy one or more capabilities. Use these examples as a general guidance for planning your environment.

Multiple Master and Worker Nodes for High Availability

In this scenario, which **deploys Intelligence with high availability**, you have three master nodes connected to three worker nodes and a database cluster. Each node runs on a separate, dedicated, connected host. All nodes have the same operating system. Each Worker Node processes events, with failover to another Worker Node if a Worker fails. All of these environments require an external server to support NFS. The Kubernetes cluster for Intelligence includes Fusion, which provides ArcSight SOAR, and Transformation Hub.


- [Diagram of this Scenario](#)
- [Characteristics of this Scenario](#)
- [Guidance for Node Configuration](#)



If this scenario resembles your intended deployment, you might want to use the `example-install-config-intelligence-high_availability.yaml` config file with the ArcSight Platform Installer. For more information about the yaml files, see [Using the Configuration Files in the Administrator's Guide for ArcSight Platform](#).

Diagram of this Scenario

Figure 2. Example deployment of Intelligence in a high-availability cluster



Characteristics of this Scenario

This scenario has the following characteristics:

- The Kubernetes cluster has three master nodes and three worker nodes, so that it can tolerate a failure of a single master and still maintain master node quorum.
- A FQDN hostname for a virtual IP is used so that clients accessing master nodes have a single reliable hostname to connect to that will shift to whatever is the current primary master node. For example, `yourdomain-ha.yourenterprise.net`.
- Transformation Hub's Kafka and ZooKeeper are deployed to all worker nodes with data replication enabled (1 original, 1 copy) so that they can tolerate a failure of a single node and still remain operational.
- Intelligence services, as well as Transformation Hub's platform and processing services, are allocated across all worker nodes so that, if one of the nodes fails, Kubernetes can move all of the components to the other node and still remain operational.
- Fusion is allocated to a single worker node.
- For the NFS configuration, use an NFS server that has high availability capabilities so that it is not a single point of failure.
- The database cluster has three nodes with data replication enabled (1 original and 1 copy) so that it can tolerate a failure of a single node and remain operational.

Guidance for Node Configuration

The worker nodes process events, with failover to another worker node in the event of a worker failure. There are no single points of failure. You need a minimum of nine physical or VM environments: three dedicated master nodes, three or more dedicated worker nodes, and a database cluster. You also need a customer-provisioned, highly available NFS server (external NFS).

The following table provides guidance for deploying the capabilities across multiple nodes to support a [large workload](#).

Node Name	Description	RAM	CPU Cores	Disk Space	Ports
<i>Master Nodes 1-3</i> masternodeNN.youenterprise.net	OMT Management Portal	256 GB	32	5 TB	OMT Vault OMT Management Portal Kubernetes NFS
<i>Database Nodes 1-3</i> databaseNN.youenterprise.net	Database	192 GB	24	28 TB	Database
<i>Worker 1</i> workernode1.youenterprise.net	Intelligence Transformation Hub	256 GB	32	5 TB	Kubernetes Transformation Hub
<i>Worker 2</i> workernode1.youenterprise.net	Intelligence Transformation Hub	256 GB	32	5 TB	Kubernetes Transformation Hub
<i>Worker 3</i> workernode1.youenterprise.net	Fusion Intelligence Transformation Hub	256 GB	32	5 TB	ArcMC Intelligence Kubernetes Transformation Hub

Single Master, Multiple Workers, and a High-availability Database

In this scenario, which **deploys Intelligence with high availability on the ArcSight Database**, you have a single master node connected to three worker nodes and a cluster for the ArcSight Database. This scenario supports an environment with modest EPS and minimal number of nodes. However, it allows for further scaling with multiple worker nodes. Each worker node runs on a separate, dedicated, connected host. All nodes have the same operating system. The Kubernetes cluster for Intelligence includes Fusion, which provides ArcSight SOAR, and Transformation Hub.


- [Diagram of this Scenario](#)
- [Characteristics of this Scenario](#)
- [Guidance for Node Configuration](#)



If this scenario resembles your intended deployment, you might want to use the example-`install-config-intelligence-scale_db.yaml` config file with the ArcSight Platform Installer. For more information about the yaml files, see [Using the Configuration Files](#) in the *Administrator's Guide for ArcSight Platform*.

Diagram of this Scenario

Figure3. Example deployment of Intelligence and Recon



Characteristics of this Scenario

This scenario has the following characteristics:

- The Kubernetes cluster overall is not highly available since it is deployed with only one master node.
- A FQDN hostname for a virtual IP is used so that clients accessing master nodes have a single reliable hostname to connect to that will shift to whatever is the current primary master node. For example, yourdomain-ha.yourenterprise.net.
- Transformation Hub's Kafka and ZooKeeper are deployed to all worker nodes with data replication enabled (1 original, 1 copy) so that they can tolerate a failure of a single node and still remain operational.
- Intelligence services, Fusion, and Transformation Hub's platform and processing services are allocated across all worker nodes so that, if one of the nodes fails, Kubernetes can move all of the components to the other node and still remain operational.
- The database cluster has three nodes with data replication enabled (1 original and 1 copy) so that it can tolerate a failure of a single node and remain operational.
- For the NFS configuration, use an NFS server that has high availability capabilities so that it is not a single point of failure.

Guidance for Node Configuration

You need a minimum of nine physical or VM environments: three dedicated master nodes, three or more dedicated worker nodes, and a database cluster. You also need a customer-provisioned, highly-available NFS server (External NFS) and an SMTP server.

The following table provides guidance for deploying the Intelligence across multiple nodes to support a [medium workload](#).

Node Name	Description	RAM	CPU Cores	Disk Space	Ports
<i>Master Node</i> masternode1.yourenterprise.net	OMT Management Portal (Optional) Fusion	256 GB	32	5 TB	OMT Management Portal Kubernetes NFS
<i>Database Nodes 1-3</i> databaseNN.yourenterprise.net	Database	192 GB	24	28 TB	Database

Node Name	Description	RAM	CPU Cores	Disk Space	Ports
<i>Worker 1</i> workernode1.youenterprise.net	Intelligence Fusion Transformation Hub	256 GB	32	5 TB	ArcMC Intelligence Kubernetes Transformation Hub
<i>Worker 2</i> workernode2.youenterprise.net	Intelligence Fusion Transformation Hub	256 GB	32	5 TB	ArcMC Intelligence Kubernetes Transformation Hub
<i>Worker 3</i> workernode3.youenterprise.net	Fusion Intelligence Transformation Hub	256 GB	32	5 TB	ArcMC Intelligence Kubernetes Transformation Hub

Everything on a Single Node—Case Management

In this scenario, which **deploys ESM Command Center on a single node**, you have the master and worker node collocated. The Kubernetes cluster for ESM Command Center includes Fusion, which provides ArcSight SOAR, and Transformation Hub.


- ["Diagram of this Scenario" on the next page](#)
- ["Characteristics of this Scenario" on the next page](#)
- ["Guidance for Node Configuration" on the next page](#)



If this scenario resembles your intended deployment, you might want to use the `example-install-config-esm_cmd_center-single-node.yaml` config file with the ArcSight Platform Installer. For more information about the yaml files, see [Using the Configuration Files](#) in the *Administrator's Guide for ArcSight Platform*. The configuration in the example file describes a single-node deployment, but you can add more worker nodes to the file.

Diagram of this Scenario

Figure 4. Example deployment of ESM Command Center on a single node



Characteristics of this Scenario

This scenario has the following characteristics:

- The Kubernetes cluster and NFS server are hosted on a single node, which creates a single point of failure. As a result, if you intend to add worker nodes, we do not recommend this configuration for high availability (HA) environments.
- FIPS 140 mode is enabled.

Guidance for Node Configuration

You need a minimum of one physical or VM environment to support master, worker, and NFS server on a single node. If you intend to install ESM Manager on the same machine, install

ESM Manager first. ESM Manager uses port 8443, so master-api-ssl-port is set to a different port to avoid a conflict.

The following table provides guidance for deploying ESM Command Center and associated capabilities on a single node to support a [small workload](#).

Node Name	Description	RAM	CPU Cores	Disk Space	Ports
Master Node yourdomain- node.yourenterprise.net	OMT Management Portal	256 GB	32	5 TB	OMT Vault OMT Management Portal Kubernetes NFS
Worker 1	ESM Command Center Fusion	32 GB	8	300 GB	ArcMC Kubernetes SOAR

Everything on a Single Node—Log Management and Compliance

In this scenario, which **deploys ArcSight Recon on a single node**, you have the master and worker node collocated with the ArcSight Database. The Kubernetes cluster for Recon includes Fusion, which provides ArcSight SOAR, and Transformation Hub.


- ["Diagram of this Scenario" on the next page](#)
- ["Characteristics of this Scenario" on the next page](#)
- ["Guidance for Node Configuration" on page 23](#)



If this scenario resembles your intended deployment, you might want to use the `example-install-config-recon-single-node.yaml` config file with the ArcSight Platform Installer. For more information about the yaml files, see [Using the Configuration Files](#) in the *Administrator's Guide for ArcSight Platform*. The configuration in the example file describes a single-node deployment, but you can add more worker nodes to the file.

Diagram of this Scenario

Figure 5. Example deployment of Recon on a single node with a database cluster



Characteristics of this Scenario

This scenario has the following characteristics:

- The Kubernetes cluster and the ArcSight Database are hosted on a single node, which creates a single point of failure. As a result, if you intend to add worker nodes, we do not recommend this configuration for high availability (HA) environments.
- FIPS 140 mode is enabled.

Guidance for Node Configuration

You need a minimum of one physical or VM environment to support master, worker, and NFS server on a single node. You can install the database cluster on the same system or a separate VM.

The following table provides guidance for deploying Recon and associated capabilities on a single node to support a [small workload](#).

Node Name	Description	RAM	CPU Cores	Disk Space	Ports
<i>Master Node</i> yourdomain- node.yourenterprise.net	OMT Management Portal	256 GB	32	5 TB	OMT Vault OMT Management Portal Kubernetes NFS
Worker 1	Recon Fusion Transformation Hub	32 GB	8	300 GB	ArcMC Kubernetes SOAR Transformation Hub
Database	Database	128 GB	24	1 TB	Database

System Performance Benchmarks for Sizing and Tuning

To help you maintain satisfactory performance of the system, we tested the system hardware under a variety of workloads in our lab environments. Based on our test results, this section identifies benchmarks that you can use to determine the system hardware required for your production environment and to tune your system's performance. These results are based on dedicated resource allocations. In virtual environments, where there is a risk of over-subscription of the physical hardware, we recommend that you ensure the system meets this sizing to avoid installation and functionality issues.

We organized the benchmarks for small, medium, and large workloads based on events per second as well as common user and system-level activities. Each day your environment might have thousands of events per second. However, the total workload depends not only on the event data received through SmartConnectors or ArcSight Enterprise Security Manager (ESM), but also other workloads that occur at the same time. For example, someone might be searching for events, new information can be coming in about the entities associated with the

events, system backup operations are performed, and a various other operations might occur. The system must be able to process all of these types of transactions simultaneously and maintain satisfactory performance.

The conditions in your environment are likely to be somewhat different than in our test lab and, as such, it is possible you might need to further adjust the system sizing or tuning values for satisfactory performance in your environment.

Options for the Database Cluster

The ArcSight Database cluster in your deployment can be collocated with the OMT cluster or in a separate, non-collocated cluster. When we tested each of these cluster options, we kept the communal storage as a separate entity. As a result, we provide dedicated sizing for the communal storage and the database node(s). We tested the following scenarios:

Collocated database cluster

Where the database, master node, and worker node are deployed on a single node as part of the OMT cluster. We used a collocated database cluster to determine the system sizing and tuning for [Small](#) and [medium](#) workloads in an on-premises, [single-node deployment](#).

Non-collocated database cluster

Where the database is deployed on dedicated nodes that make up the database cluster, and this cluster is not a part of the OMT cluster. We used a non-collocated database cluster to determine the system sizing and tuning for the following scenarios:

- [Large](#) workload in an on-premises, [multi-node deployment](#)
- [Small](#), [medium](#), and large workloads in a cloud, multi-node deployment

SOAR Focused Workload

ArcSight SOAR is embedded in the Fusion capability. This section describes the system sizing and tuning for the ArcSight Platform and the deployed capabilities of Fusion and (optionally) Command Center for ESM. These capabilities are integrated with a full ESM deployment that is deployed separately.

In this scenario, since only the Fusion and Command Center for ESM capabilities are deployed on the ArcSight Platform and these capabilities are handling a small workload (as they are not in the events processing pipeline), it is possible to deploy this on a single machine.

System Sizing

This section describes the system sizing of the tested system for a single machine.

On-Premises Deployment

The *OMT Master/Worker Node* system resources are where the core platform and capabilities are deployed in an all-in-one co-located configuration on the tested system.

Category	1 x OMT Master/Worker Node
Processor	Intel(R) Xeon(R) Gold 6248 CPU @ 2.50GHz
vCPU(s) (# threads)	24
RAM (per node)	128 GB
Disks (per node)	1
Total disk space (per node)	750 GB
SOAR DB storage on NFS arcsight-volume	1 TB (holds up to 1000 case related events per day for two years)
K-safety level	0

Small Workload



As we complete testing for additional scenarios, we will add more information to this page.

This section describes the system sizing and tuning results from tests of the ArcSight Platform and deployed capabilities Transformation Hub, Fusion, Command Center for ESM, Intelligence, Recon, and the ArcSight Database that has been confirmed in our testing lab to maintain satisfactory performance of the system under a small workload.

- [Workloads](#)
- [System Sizing](#)
- [System Tuning](#)

Workloads

This section describes the workload that was placed on the tested system.

- [Event Workload](#)
- [Other Workload](#)

Event Workload

Application	Events per second
Microsoft Windows	700
InfoBlox NIOS	700
Intelligence Data (VPN, AD, Proxy)	100
Total	1,500

Other workload

Category	Level
Storage Groups	10
Searches	3 per hour (concurrent)
Reports	1 scheduled every hour

System Sizing

This section describes the sizing of the tested system.

- [On-Premises Deployment](#)
- [AWS Deployment](#)
- [Azure Deployment](#)

On-Premises Deployment

The *OMT Master/Worker Node/Database Node* system resources are where the core platform, Transformation Hub, Fusion, Command Center for ESM, Recon, and Database compute components were deployed in an all-in-one collocated configuration on the tested system. However, the Database *Communal Storage* components were deployed on a separate node because they are not embedded within the ArcSight Platform. When using this information as guidance for your own system sizing, the OMT Master/Worker Node/Database Node system resources are always needed, but the Database Communal Storage system resources are only needed when deploying Recon or Intelligence.

Category	1 x OMT Master/Worker Node/Database Node	1 x Communal Storage
Processor	Intel(R) Xeon(R) Gold 6248 CPU @ 2.50GHz	Intel(R) Xeon(R) Gold 6248 CPU @ 2.50GHz

vCPU(s) (# threads)	24	6
RAM (per node)	128 GB	32 GB
Disks (per node)	ESX data store	ESX data store
Storage per day (1x)	7 GB (depot) + 15 GB (ES)	27 GB (MinIO)
Total disk space (5 Billion events)	1 TB (holds up to 45 days of events)	1 TB (holds up to 40 days of events)
K-safety level	0	N/A

AWS Deployment

The *OMT Worker (Platform)* system resources are where the core platform, Transformation Hub, Fusion, Command Center for ESM, and Recon components were deployed on the tested system. However, Intelligence components were deployed on the *OMT Worker (Intelligence)* system resources because they utilize a significant amount of resources when running analytics jobs. When using this information as guidance for your own system sizing, the *OMT Worker (Platform)* system resources are always needed, the *Database* system resources are only needed when deploying Recon or Intelligence, and the *OMT Worker (Intelligence)* system resources are only needed when deploying ArcSight Intelligence.

Category	OMT Worker (Platform)	Database	OMT Worker (Intelligence)
Instance Type	m5.2xlarge	m5d.4xlarge	m5.2xlarge
Instance Count	3	3	3
Disks (per node)	500 GB - EBS storage (gp2)	2 x 300 NVMe SSD	500 GB - EBS storage (gp2)

Azure Deployment

The *OMT Worker (Platform)* system resources are where the core platform, Transformation Hub, Fusion, Command Center for ESM, and Recon components were deployed on the tested system. However, Intelligence components were deployed on the *OMT Worker (Intelligence)* system resources because they utilize a significant amount of resources when running analytics jobs. When using this information as guidance for your own system sizing, the "OMT Worker (Platform)" system resources are always needed, the *Database* system resources are only needed when deploying Recon or Intelligence, and the *OMT Worker (Intelligence)* system resources are only needed when deploying ArcSight Intelligence.

Category	OMT Worker (Platform)	Database	OMT Worker (Intelligence)
Instance Type	D2s_V3	D4s_V3	D2s_V3

Instance Count	3	3	3
Disks (per node)	1 x 500 GB - Premium SSD	2 x 300 GB - Premium SSD	1 x 500 GB - Premium SSD

System Tuning

This section describes the system tuning of the tested system.

- [Database Tuning](#)
- [Transformation Hub Tuning](#)
- [Intelligence Tuning](#)
- [Fusion Tuning](#)
- [SmartConnector Tuning](#)

Database Tuning

Category	Property	On-Premises	AWS	Azure
Core Database	shard_count	3	3	3
Core Database	depot_size	40%	60%	60%
Tuple Mover	tm_concurrency	5	6	5
Tuple Mover	tm_memory	10G	10G	10G
Tuple Mover	plannedconcurrency	5	6	5
Tuple Mover	tm_memory_usage	10000	10000	10000
Tuple Mover	maxconcurrency	10	7	10
Ingest Resource pools	ingest_pool_memory_size	30%	30%	30%
Ingest Resource pools	ingest_pool_planned_concurrency	6	6	6
Backup	Backup Interval (hours)	1	1	1
Communal Storage	Server-side Encryption	Disabled	Disabled	Yes (MMK)

Transformation Hub Tuning

Property	On-Premises	AWS	Azure
# of Kafka broker nodes in the Kafka cluster	1	3	3

# of ZooKeeper nodes in the ZooKeeper cluster	1	3	3
# of Partitions assigned to each Kafka Topic*	12	24	24
# of replicas assigned to each Kafka Topic	1	2	2
# of message replicas for the __consumer_offsets Topic	1	3	3
Schema Registry nodes in the cluster	1	3	3
# of CEF-to-Avro Stream Processor instances to start**	0	0	0
# of Enrichment Stream Processor Group instances to start	2	2	2

*Kafka topics - th-arcsight-avro; mf-event-avro-enriched; and th-cef, if connectors are configured to send to Transformation Hub in CEF format

**If connectors are configured to send Avro format to Transformation Hub, you can set the # of CEF-to-Avro Stream Processor instances to start quantity to 0 because there is no need to convert CEF to Avro.

Intelligence Tuning

Property	On-Premises	AWS	Azure
Elasticsearch Shard Count	6	6	6
Elasticsearch data processing Instances	1	3	3
Elasticsearch Index Replica Count	0	1	1
Elasticsearch Memory (GB)	10	4	4
Elasticsearch number of cores	6	2	2
Elasticsearch Size Per Batch	5mb	5mb	5mb
Logstash Instances	2	3	3
Logstash pipeline workers per instance	2	1	1
Logstash Pipeline Batch size	500	500	500
LogStash Filter Applied	yes	yes	yes
Spark parallelism	32	32	32
Spark number of executors	3	3	3
Spark executor memory	5g	4g	4g
Spark number of executor cores	1	1	1
Spark driver memory	4g	4g	3g
Spark memory overhead factor	0.2	0.2	0.2
Intelligence Job per day	1	1	1

Fusion Tuning

Category	All Deployments
Event Integrity Check Task Count	1
Event Integrity Check Chunk Size	1000
Use Event Integrity Check Resource Pool	false

SmartConnector Tuning

Category	All Deployments
SmartConnector version that we tested	8.3.0.14008.0
Instance Count	1
Acknowledgement Mode	leader
usessl (Transformation Hub Destination Param)	false
contenttype (Transformation Hub Destination Param)	Avro
topic (Transformation Hub Destination Param)	th-arcSight-avro
compression.type	gzip
transport.batchqueuesize	20000
transport.cefkafka.batch.size	50000
transport.cefkafka.linger.ms	10
transport.cefkafka.max.request.size	4194304
transport.cefkafka.multiplekafkaProducers	true
transport.cefkafka.threads	3

Medium Workload



As we complete testing for additional scenarios, we will add more information to this page.

This section describes the system sizing and tuning results from tests of the ArcSight Platform and deployed capabilities Transformation Hub, Fusion, Command Center for ESM, Intelligence, Recon, and the ArcSight Database that has been confirmed in our testing lab to maintain satisfactory performance of the system under a medium workload.

- [Workloads](#)
- [System Sizing](#)
- [System Tuning](#)

Workloads

This section describes the workload that was placed on the tested system.

- [Event Workload](#)
- [Other Workload](#)

Event Workload

This table provides event ingestion workload in events per second:

Application	On-Premises Collocated Database	AWS or Azure Non-collocated Database
Microsoft Windows	2,400	9,000
InfoBlox NIOS	2,400	9,000
Intelligence Data (VPN, AD, Proxy)	250	2,000
Total	5,050	20,000

Other Workload

Category	Level
Storage Groups	10
Searches	3 per hour (concurrent)
Reports	1 scheduled every hour

System Sizing

This section describes the system sizing of the tested system.

- [On-Premises Deployment](#)
- [AWS Deployment](#)
- [Azure Deployment](#)

On-Premises Deployment

The *OMT Master/Worker Node/Database Node* system resources are where the core platform, Transformation Hub, Fusion, Command Center for ESM, Recon, and Database compute components were deployed in an all-in-one collocated configuration on the tested system. However, the Database *Communal Storage* components were deployed on a separate node because they are not embedded within the ArcSight Platform. When using this information as guidance for your own system sizing, the OMT Master/Worker Node/Database Node system resources are always needed, but the Database Communal Storage system resources are only needed when deploying Recon or Intelligence.

Category	1 x OMT Master/Worker Node/Database Node	1 x Communal Storage
Processor	Intel(R) Xeon(R) Gold 6248 CPU @ 2.50GHz	Intel(R) Xeon(R) Gold 6248 CPU @ 2.50GHz
vCPU(s) (# threads)	48	8
RAM (per node)	192 GB	48 GB
Disks (per node)	ESX data store	ESX data store
Storage per day (1x)	10 GB (depot) + 20 GB (ES)	100 GB (MinIO)
Total disk space (5 Billion events)	1 TB (holds up to 30 days of events)	1 TB (holds up to 15 days of events)
K-safety level	0	N/A

AWS Deployment

The *OMT Worker (Platform)* system resources are where the core platform, Transformation Hub, Fusion, Command Center for ESM, and Recon components were deployed on the tested system. However, Intelligence components were deployed on the *OMT Worker (Intelligence)* system resources because they utilize a significant amount of resources when running analytics jobs. When using this information as guidance for your own system sizing, the OMT Worker (Platform) system resources are always needed, the *Database* system resources are only needed when deploying Recon or Intelligence, and the OMT Worker (Intelligence) system resources are only needed when deploying ArcSight Intelligence.

Category	OMT Worker (Platform)	Database	OMT Worker (Intelligence)
Instance Type	m5.4xlarge	m5.12xlarge	m5.4xlarge
Instance Count	3	6	3
Disks (per node)	1 X 2048 GB (gp3) EBS Volumes	8 x 250 GB (gp3) EBS Volumes	1 X 2048 GB (gp3) EBS Volumes

Azure Deployment

The *OMT Worker (Platform)* system resources are where the core platform, Transformation Hub, Fusion, Command Center for ESM, and Recon components were deployed on the tested system. However, Intelligence components were deployed on the *OMT Worker (Intelligence)* system resources because they utilize a significant amount of resources when running analytics jobs. When using this information as guidance for your own system sizing, the "OMT Worker (Platform)" system resources are always needed, the *Database* system resources are only needed when deploying Recon or Intelligence, and the OMT Worker (Intelligence) system resources are only needed when deploying ArcSight Intelligence.

Category	OMT Worker (Platform)	Database	OMT Worker (Intelligence)
Instance Type	D16s_V3	D32s_V3	D16s_V3
Instance Count	3	6	3
Disks (per node)	2 TB - Premium SSD	2 TB (Depot) - Premium SSD	2 TB - Premium SSD

System Tuning

This section describes the system tuning of the tested system.

- [Database Tuning](#)
- [Transformation Hub Tuning](#)
- [Intelligence Tuning](#)
- [Fusion Tuning](#)
- [SmartConnector Tuning](#)

Database Tuning

Category	Property	On-Premises	Azure	AWS
Core Database	shard_count	3	18	18
Core Database	depot_size	40%	60%	60%
Tuple Mover	tm_concurrency	5	5	10
Tuple Mover	tm_memory	10G	10G	10G
Tuple Mover	plannedconcurrency	5	5	5
Tuple Mover	tm_memory_usage	10000	10000	20000
Tuple Mover	maxconcurrency	10	10	10

Ingest Resource pools	ingest_pool_memory_size	30%	30%	30%
Ingest Resource pools	ingest_pool_planned_concurrency	6	6	6
Backup	Backup Interval (hours)	1	1	1

Transformation Hub Tuning

Property	On-Premises	Azure	AWS
# of Kafka broker nodes in the Kafka cluster	1	3	3
# of ZooKeeper nodes in the ZooKeeper cluster	1	3	3
# of Partitions assigned to each Kafka Topic*	12	72	72
# of replicas assigned to each Kafka Topic	1	2	2
# of message replicas for the __consumer_offsets Topic	1	3	3
Schema Registry nodes in the cluster	1	3	3
# of CEF-to-Avro Stream Processor instances to start**	0	0	3
# of Enrichment Stream Processor Group instances to start	2	3	3

*Kafka topics - th-arcsight-avro; mf-event-avro-enriched; and th-cef, if connectors are configured to send to Transformation Hub in CEF format

**If connectors are configured to send Avro format to Transformation Hub, you can set the # of CEF-to-Avro Stream Processor instances to start quantity to 0 because there is no need to convert CEF to Avro.

Intelligence Tuning

Property	On-Premises	Azure	AWS
Elasticsearch Shard Count	6	6	6
Elasticsearch data processing Instances	1	3	3
Elasticsearch Index Replica Count	0	1	1
Elasticsearch Memory (GB)	14	12	12
Elasticsearch number of cores	8	6	5
Elasticsearch Size Per Batch	5mb	5mb	5mb
Logstash Instances	3	12	15
Logstash pipeline workers per instance	2	2	1
Logstash Pipeline Batch size	500	1000	500
LogStash Filter Applied	yes	yes	yes
Spark Parallelism	32	32	64

Spark number of executors	3	8	9
Spark executor memory	6g	8g	7g
Spark number of executor cores	1	1	1
Spark Driver Memory	6g	8g	8g
Spark Memory Overhead Factor	0.2	0.2	0.2
Intelligence Job per day	1	1	1

Fusion Tuning

Category	All Deployments
Event Integrity Check Task Count	1
Event Integrity Check Chunk Size	1000
Use Event Integrity Check Resource Pool	false

SmartConnector Tuning

Category	All Deployments
SmartConnector version that we tested	8.3.0.14008.0
Instance Count	1
Acknowledgement Mode	none
usessl (Transformation Hub Destination Param)	false
contenttype (Transformation Hub Destination Param)	Avro
topic (Transformation Hub Destination Param)	th-arcsight-avro
compression.type	gzip
transport.batchqueuesize	20000
transport.cefkafka.batch.size	50000
transport.cefkafka.linger.ms	10
transport.cefkafka.max.request.size	4194304
transport.cefkafka.multiplekafkaproducers	true
transport.cefkafka.threads	6
syslog.parser.threadcount	6

Large Workload



As we complete testing for additional scenarios, we will add more information to this page.

This section describes the system sizing and tuning results from tests of the ArcSight Platform and deployed capabilities Transformation Hub, Fusion, Command Center for ESM, Intelligence, Recon, and the ArcSight Database that has been confirmed in our testing lab to maintain satisfactory performance of the system under a large workload.

- [Workloads](#)
- [System Sizing](#)
- [System Tuning](#)

Workloads

This section describes the workload that was placed on the tested system.

- [Event Workload](#)
- [Other Workload](#)

Event Workload

Application	Events per second
Microsoft Windows	54,000
InfoBlox NIOS	54,000
Intelligence Data (VPN, AD, Proxy)	12,000
Total	120,000

Other Workload

Category	Level
Storage Groups	10
Searches	3 per hour (concurrent)
Reports	1 scheduled every hour

System Sizing

This section describes the system sizing of the tested system.

AWS Deployment

The *OMT Worker (Platform)* system resources are where the core platform, Transformation Hub, Fusion, Command Center for ESM, and Recon components were deployed on the tested system. However, Intelligence components were deployed on the *OMT Worker (Intelligence)* system resources because they utilize a significant amount of resources when running analytics jobs. When using this information as guidance for your own system sizing, the OMT Worker (Platform) system resources are always needed, the *Database* system resources are only needed when deploying Recon or Intelligence, and the OMT Worker (Intelligence) system resources are only needed when deploying ArcSight Intelligence.

Category	OMT Worker (Platform)	Database	OMT Worker (Intelligence)
Instance Type	m5.8xlarge	m5d.8xlarge	m5.8xlarge
Instance Count	3	18	6
Disks (per node)	3 TB - EBS storage	2 x 600 NVMe SSD	3 TB - EBS storage

System Tuning

This section describes the system tuning of the tested system.

- [Database Tuning](#)
- [Transformation Hub Tuning](#)
- [Intelligence Tuning](#)
- [Fusion Tuning](#)
- [SmartConnector Tuning](#)

Database Tuning:

Category	Property	AWS
Core Database	shard_count	18
Core Database	depot_size	60%
Tuple Mover	tm_concurrency	5
Tuple Mover	tm_memory	10G
Tuple Mover	plannedconcurrency	5
Tuple Mover	tm_memory_usage	10000
Tuple Mover	maxconcurrency	10

Ingest Resource pools	ingest_pool_memory_size	30%
Ingest Resource pools	ingest_pool_planned_concurrency	6
Backup	Backup Interval (hours)	1
Communal Storage	Server-side Encryption	disabled

Transformation Hub Tuning

Property	AWS
# of Kafka broker nodes in the Kafka cluster	3
# of ZooKeeper nodes in the ZooKeeper cluster	3
# of Partitions assigned to each Kafka Topic*	108
# of replicas assigned to each Kafka Topic	2
# of message replicas for the __consumer_offsets Topic	3
Schema Registry nodes in the cluster	3
# of CEF-to-Avro Stream Processor instances to start**	0
# of Enrichment Stream Processor Group instances to start	6

*Kafka topics - th-arcsight-avro; mf-event-avro-enriched; and th-cef, if connectors are configured to send to Transformation Hub in CEF format

**If connectors are configured to send Avro format to Transformation Hub, you can set the # of CEF-to-Avro Stream Processor instances to start quantity to 0 because there is no need to convert CEF to Avro.

Kafka Override Parameters	AWS
arcsight.eventbroker.kafka.KAFKA_NUM_IO_THREADS	256
arcsight.eventbroker.kafka.KAFKA_NUM_NETWORK_THREADS	52
arcsight.eventbroker.kafka.KAFKA_NUM_REPLICA_FETCHERS	145

Intelligence Tuning

Property	AWS
Elasticsearch Shard Count	6
Elasticsearch data processing Instances	6
Elasticsearch Index Replica Count	1
Elasticsearch Memory (GB)	24
Elasticsearch number of cores	12

Elasticsearch Size Per Batch	10mb
Logstash Instances	108
Logstash pipeline workers per instance	2
Logstash Pipeline Batch size	2000
Spark Parallelism	64
Spark number of executors	24
Spark executor memory	12g
Spark number of executor cores	1
Spark Driver Memory	8g
Spark Memory Overhead Factor	0.2
Intelligence Job per day	1

Fusion Tuning

Category	All Deployments
Event Integrity Check Task Count	6
Event Integrity Check Chunk Size	1000
Use Event Integrity Check Resource Pool	false

SmartConnector Tuning

Category	All Deployments
SmartConnector version that we tested	8.3.0.14008.0
Instance Count	5
Acknowledgement Mode	none
usessl (Transformation Hub Destination Param)	false
contenttype (Transformation Hub Destination Param)	Avro
topic (Transformation Hub Destination Param)	th-arcsight-avro
compression.type	gzip
transport.batchqueuesize	20000
transport.cefkafka.batch.size	50000
transport.cefkafka.linger.ms	10

transport.cefkafka.max.request.size	4194304
transport.cefkafka.multiplekafkaproducers	true
transport.cefkafka.threads	6

Publication Status

Released: July 27, 2023

Send Documentation Feedback

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

Feedback on ArcSight Platform Technical Requirements (ArcSight Platform 23.2)

Just add your feedback to the email and click send.

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

We appreciate your feedback!

Legal Notices

Copyright Notice

© Copyright 2001 - 2023 Micro Focus or one of its affiliates

Confidential computer software. Valid license from Micro Focus required for possession, use or copying. The information contained herein is subject to change without notice.

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

No portion of this product's documentation may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's internal use, without the express written permission of Micro Focus.

Notwithstanding anything to the contrary in your license agreement for Micro Focus ArcSight software, you may reverse engineer and modify certain open source components of the software in accordance with the license terms for those particular components. See below for the applicable terms.

U.S. Governmental Rights. For purposes of your license to Micro Focus ArcSight software, "commercial computer software" is defined at FAR 2.101. If acquired by or on behalf of a civilian agency, the U.S. Government acquires this commercial computer software and/or commercial computer software documentation and other technical data subject to the terms of the Agreement as specified in 48 C.F.R. 12.212 (Computer Software) and 12.211 (Technical Data) of the Federal Acquisition Regulation ("FAR") and its successors. If acquired by or on behalf of any agency within the Department of Defense ("DOD"), the U.S. Government acquires this commercial computer software and/or commercial computer software documentation subject to the terms of the Agreement as specified in 48 C.F.R. 227.7202-3 of the DOD FAR Supplement ("DFARS") and its successors. This U.S. Government Rights Section 18.11 is in lieu of, and supersedes, any other FAR, DFARS, or other clause or provision that addresses government rights in computer software or technical data.

For information about legal notices, trademarks, disclaimers, warranties, export and other use restrictions, U.S. Government rights, patent policy, and FIPS compliance, see <https://www.microfocus.com/about/legal/>.