Identity Intelligence 1.1
System Requirements

April 2020

NetIQ

Legal Notice

About These System Requirements

Micro Focus recommends the tested platforms listed below. However, customers running on any platforms not provided in this list or with untested configurations will be supported until the point Micro Focus determines that the root cause is the untested platform or configuration. Issues that can be reproduced on the tested platforms will be prioritized and fixed according to standard defect-handling policies.

- Chapter 1, “Software Requirements,” on page 7
- Chapter 2, “Hardware Requirements and Tuning Guidelines,” on page 11
- Chapter 3, “Network File System,” on page 21
- Chapter 4, “Ports Used,” on page 23

For more information about support polices, see Support Policies.

For information about installation, see the *Administrator Guide for Identity Intelligence*.

Additional Documentation

The Identity Intelligence documentation library includes the following resources:

- *Administrator Guide to Identity Intelligence*, which provides information about deploying, configuring, and maintaining this product
- *User Guide to Identity Intelligence*, which is embedded in the product to provide both contextual Help and conceptual information
- *Release Notes for Identity Intelligence*

For the most recent version of the system requirements and other Identity Intelligence documentation resources, visit the documentation for Identity Intelligence.

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/.

Contents

1 Software Requirements

This section lists the software needed to install and run Identity Intelligence.

Software Requirements for Identity Intelligence 1.1.2

Category	Minimum Requirement
Operating systems	A **minimal** installation of Red Hat Enterprise Linux 7.7 (x86, x64)
File systems	One of the following: ◆ EXT3 ◆ EXT4 (recommended) ◆ Logical Volume Manager (LVM) ◆ XFS
Container Deployment Foundation (CDF)	CDF 2020.05 **NOTE:** If you do not already have CDF, Identity Intelligence includes the installation files.
Database	Vertica 9.2.1 **NOTE:** If you do not already have database in your environment, Identity Intelligence includes the installation files.
Data Processing	Transformation Hub 3.3 **NOTE:** If you do not already have Transformation Hub in your environment, Identity Intelligence includes the installation images for deployment. For documentation, see Transformation Hub documents.
Fusion	Fusion 1.1.0 **NOTE:** If you do not already have Fusion in your environment, Identity Intelligence includes the installation images.
Data Collection	Data Collection Software
Browser	◆ Google Chrome ◆ Mozilla Firefox **NOTE:** Browsers should not use a proxy to access Container Deployment Foundation (CDF) application because this might result in inaccessible web pages.

Data Collection Software

Identity Intelligence works with the following data collection software:

Product or Component	Version	Notes
Identity Manager	4.7.4 or later	
Identity Manager Driver for Entity Data Model	1.1 or later	*Required for use with Identity Manager* Provided with the Identity Intelligence download For more information about the system requirements for Identity Manager Driver for Entity Data Model, see System Requirements.
Identity Governance	3.5.2 or later	
IG Configuration Utility for Entity Data Model	1.0 or later	*Required for use with Identity Governance* Provided with the Identity Intelligence download
SmartConnector for Syslog NG Daemon	7.14 or later	Provided with the Identity Intelligence download

Software Requirements for Identity Intelligence 1.1.1 and 1.1

Category	Minimum Requirement
Operating systems	A **minimal** installation of Red Hat Enterprise Linux 7.7 (x86, x64)
File systems	One of the following: ◆ EXT3 ◆ EXT4 (recommended) ◆ Logical Volume Manager (LVM) ◆ XFS
Container Deployment Foundation (CDF)	CDF 2020.02 **NOTE:** If you do not already have CDF, Identity Intelligence includes the installation files.
Database	Vertica 9.2.1 **NOTE:** If you do not already have database in your environment, Identity Intelligence includes the installation files.
Data Processing	Transformation Hub 3.2 **NOTE:** If you do not already have Transformation Hub in your environment, Identity Intelligence includes the installation images for deployment. For documentation, see Transformation Hub documents.

Category	Minimum Requirement
Analytics	Analytics 3.1.0 **NOTE:** If you do not already have Analytics in your environment, Identity Intelligence includes the installation images.
Data Collection	Data Collection Software
Browser	◆ Google Chrome ◆ Mozilla Firefox **NOTE:** Browsers should not use a proxy to access Container Deployment Foundation (CDF) application because this might result in inaccessible web pages.

Data Collection Software

Identity Intelligence works with the following data collection software:

Product or Component	Version	Notes
Identity Manager	4.7.4 or later	
Identity Manager Driver for Entity Data Model	1.1 or later	*Required for use with Identity Manager* Provided with the Identity Intelligence download For more information about the system requirements for Identity Manager Driver for Entity Data Model, see System Requirements.
Identity Governance	3.5.2 or later	
IG Configuration Utility for Entity Data Model	1.0 or later	*Required for use with Identity Governance* Provided with the Identity Intelligence download
SmartConnector for Syslog NG Daemon	7.14 or later	Provided with the Identity Intelligence download

2 Hardware Requirements and Tuning Guidelines

The guidelines in this section are for a deployment where you install all of the following software:

- Database
- Transformation Hub
- Identity Intelligence

The hardware requirements are based on dedicated resource allocations. In virtual environments, where there is a risk of over subscription of the physical hardware, ensure that the Identity Intelligence system meets these hardware requirements to avoid installation and functionality issues.

- "Understanding the Workload for Identity Intelligence" on page 11
- "System Sizing for Single-Node Deployment" on page 11
- "System Sizing for High Availability" on page 17
- "Disk Partition Requirements" on page 19

NOTE: The system sizing was tested in an Identity Intelligence environment without SSL communication.

Understanding the Workload for Identity Intelligence

The total workload for Identity Intelligence depends on your data sources, such as Identity Manager, and the number of identity governance transactions that occur within those data sources each day. For example, each day, your environment might have thousands of identities requesting one of the thousands of access rights. Or, identities approving and denying those requests. At the same time, someone might be modifying the properties of those identities and entities. Identity Intelligence must be able to process all of these types of transactions. Thus, this document lists requirements for small, medium, and large workloads.

System Sizing for Single-Node Deployment

- "System Sizing for a Small Workload" on page 12
- "System Sizing for Medium Workload" on page 13
- "System Sizing for Large Workload" on page 15

System Sizing for a Small Workload

This section helps you in determining whether your environment might meet the requirements for a small workload environment. It provides guidance for hardware requirements and tuning the performance of the workload. You might compare this information with the guidance for medium and large workloads.

Workload Distribution for a Small Workload

The following table provides an example of how identity governance activities might occur in a small workload:

Application	Category	Expected Workload
Identity Governance	Identities	5,000
	Deltas (per day)	10,000
Identity Manager	Identities	5,000
	Events per second	10
Identity Intelligence (total workload)	Identities	10,000
	Transactions (per day)	20,000
	Views (concurrent)	5

System Sizing for a Small Workload

Category	Requirement
Cluster node (master+worker)	1
CPU cores (per node)	8
RAM (per node)	32
Disks (per node)	2
Storage per day (1x)	350 MB
Total disk space (365 days)	500 GB

Database Resource Pools Tuning for a Small Workload

Category	Property	Value
Database	active_partitions	3
	tm_concurrency	4
	tm_memory	2,000
Resource pools	ingest_pool_memory_size	10%
	mf_entity_ingest_pool_memory_size	10%
	mf_entity_ingest_pool_planned_concurrency	4
Scheduler	plannedconcurrency	3
	tm_memory_usage	2,000
	maxconcurrency	4

Transformation Hub Tuning for a Small Workload

Property	Quantity
# of Kafka broker nodes in the Kafka cluster	1
# of ZooKeeper nodes in the ZooKeeper cluster	1
# of Partitions assigned to each Kafka Topic	1
# of replicas assigned to each Kafka Topic	1
# of message replicas for the __consumer_offsets Topic	1
Schema Registry nodes in the cluster	1
Kafka nodes required to run Schema Registry	1
# of CEF-to-Avro Stream Processor instances to start	1

System Sizing for Medium Workload

This section helps you in determining whether your environment meets the requirements for a medium workload environment. It provides guidance for hardware requirements and tuning the performance of the workload. You might compare this information with the guidance for small and large workloads.

- "Workload Distribution for a Medium Workload" on page 14
- "System Sizing for a Medium Workload" on page 14

Workload Distribution for a Medium Workload

The following table provides an example of how identity governance activities might occur in a medium workload:

Application	Category	Expected Workload
Identity Governance	Identities	25,000
	Accounts	50,000
	Groups	7,500
	Entitlements / permissions	7,500
	Deltas (per day)	200,000
Identity Manager	Identities	25,000
	Events per second	12
Identity Intelligence (total workload)	Identities	50,000
	Transactions (per day)	100,000
	Views (concurrent)	5

System Sizing for a Medium Workload

Category	Requirement
Cluster node (master+worker)	1
CPU cores (per node)	8
RAM (per node)	48
Disks (per node)	2
Storage per day (1x)	950 MB
Total disk space (365 days)	1 TB

Database Resource Pools Tuning for a Medium Workload

Category	Property	Value
Database	active_partitions	4
	tm_concurrency	5
	tm_memory	4,000
Resource pools	ingest_pool_memory_size	10%
	mf_entity_ingest_pool_memory_size	30%
	mf_entity_ingest_pool_planned_concurrency	12
Scheduler	plannedconcurrency	4
	tm_memory_usage	4,000
	maxconcurrency	5

Transformation Hub Tuning for a Medium Workload

Property	Quantity
# of Kafka broker nodes in the Kafka cluster	1
# of ZooKeeper nodes in the ZooKeeper cluster	1
# of Partitions assigned to each Kafka Topic	3
# of replicas assigned to each Kafka Topic	1
# of message replicas for the __consumer_offsets Topic	1
Schema Registry nodes in the cluster	1
Kafka nodes required to run Schema Registry	1
# of CEF-to-Avro Stream Processor instances to start	1

System Sizing for Large Workload

This section helps you in determining whether your environment might meets the requirements for a large workload environment. It provides guidance for hardware requirements and tuning the performance of the workload. You might compare this information with the guidance for small and medium workloads.

- "Workload Distribution for a Large Workload" on page 16
- "System Sizing for a Large Workload" on page 16

- “Database Resource Pools Tuning for a Large Workload” on page 17
- “Transformation Hub Tuning for a Large Workload” on page 17

Workload Distribution for a Large Workload

The following table provides an example of how identity governance activities might occur in a large workload:

Application	Category	Expected Workload
Identity Governance	Identities	500,000
	Accounts	500,000
	Groups	7,500
	Entitlements / permissions	7,500
	Deltas (per day)	500,000
Identity Manager	Identities	500,000
	Events per second	100
Identity Intelligence (total workload)	Identities	1 million
	Transactions (per day)	1 million
	Views (concurrent)	10

System Sizing for a Large Workload

Category	Requirement
Cluster node (master+worker)	1
CPU cores (per node)	12
RAM (per node)	64
Disks (per node)	2
Storage per day (1x)	10 GB
Total disk space (365 days)	3 TB

Database Resource Pools Tuning for a Large Workload

Category	Property	Value
Database	active_partitions	6
	tm_concurrency	7
	tm_memory	6,000
Resource pools	ingest_pool_memory_size	15%
	mf_entity_ingest_pool_memory_size	30%
	mf_entity_ingest_pool_planned_concurrency	24
Scheduler	plannedconcurrency	6
	tm_memory_usage	6,000
	maxconcurrency	7

Transformation Hub Tuning for a Large Workload

Property	Quantity
# of Kafka broker nodes in the Kafka cluster	1
# of ZooKeeper nodes in the ZooKeeper cluster	1
# of Partitions assigned to each Kafka Topic	6
# of replicas assigned to each Kafka Topic	1
# of message replicas for the __consumer_offsets Topic	1
Schema Registry nodes in the cluster	1
Kafka nodes required to run Schema Registry	1
# of CEF-to-Avro Stream Processor instances to start	1

System Sizing for High Availability

This section helps you in determining whether your environment might meets the requirements for a high availability environment. It provides guidance for hardware requirements and tuning the performance.

System Sizing for High Availability

Master Node

Category	Requirement
Master nodes	3
CPU cores (per node)	4
RAM (per node)	16 GB
Disks (per node)	2
Hard disk	350 GB

Worker Node

Category	Requirement for Small Workloads	Requirement for Medium Workloads	Requirement for Large Workloads
Worker nodes	2	2	2
CPU cores (per node)	8	8	12
RAM (per node)	32 GB	48 GB	64 GB
Disks (per node)	2	2	2
Storage per day (1x)	350 MB	950 MB	10 GB
Total disk space (365 days)	500 GB	1 TB	3 TB

External NFS Server

Category	Requirement
CPU cores	4
RAM	4 GB
Hard disk	300 GB

Database Resource Pools Tuning for High Availability

Based on the workload, see the relevant (small, medium, large) workload distribution section.

Transformation Hub Tuning for High Availability

Property	Quantity
# of Kafka broker nodes in the Kafka cluster	3
# of ZooKeeper nodes in the ZooKeeper cluster	3
# of Partitions assigned to each Kafka Topic	6
# of replicas assigned to each Kafka Topic	2
# of message replicas for the __consumer_offsets Topic	3
Schema Registry nodes in the cluster	3
Kafka nodes required to run Schema Registry	3
# of CEF-to-Avro Stream Processor instances to start	1

Disk Partition Requirements

This section lists the minimum disk space needed to run Identity Intelligence. In some deployments, you might deploy Identity Intelligence with Recon.

- "Disk Partition Requirements for Master Node" on page 19
- "Disk Partition Requirements for Worker Node" on page 20

Disk Partition Requirements for Master Node

Partition	Requirement
`/`	50 GB
`/boot`	250 MB
`/opt`	250 GB
`swap`	16 GB

Disk Partition Requirements for Worker Node

Partition	Small Workload	Medium Workload	Large Workload	Notes
/	50 GB	50 GB	50 GB	Contains operating system and SmartConnector files.
/boot	250 MB	250 MB	250 MB	Boot partition
/opt	250 GB	350 GB	450 GB	Contains Transformation Hub (Kafka) and NFS files.
/opt/vertica/data	250 GB	500 GB	2.5 TB	Contains database which stores the data collected from data sources.
swap	16 GB	16 GB	16 GB	

NOTE: Instead of /opt/vertica/data, you can specify any other directory that has adequate space to store data during database installation.

3 Network File System

Identity Intelligence supports several options for a network file system (NFS).

Required File Systems

Category	Minimum Requirement
NFS Types	◆ Amazon EFS ◆ HPE 3PAR File Persona ◆ Linux-based NFS ◆ NetApp
NFS Server Versions	◆ NFSv4 ◆ NFSv3

Minimum Directory Sizes for the NFS

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

Directory	Minimum Size
{NFS_ROOT_DIRECTORY}/itom/itom-vol	130 GB
{NFS_ROOT_DIRECTORY}/itom/db-single-vol	Depends, but start with 10 GB
{NFS_ROOT_DIRECTORY}/itom/db-backup-vol	Depends, but start with 10 GB
{NFS_ROOT_DIRECTORY}/itom/itom-logging-vol	Depends, but start with 40 GB
{NFS_ROOT_DIRECTORY}/arcsight-vol	10 GB

4 Ports Used

Identity Intelligence uses following firewall ports. Therefore, ensure that the following ports are available.

Database

We do not recommend placing a firewall between nodes (all nodes should be behind a firewall), but if you must use a firewall between nodes, ensure the following ports are available.

For more information about database ports, see Vertica Documentation.

Ports	Protocol	Node	Description
22	TCP	All database nodes	Used by the Administration Tools and the Management Console Cluster installation wizard.
5433	TCP	All database nodes	Used by database clients, such as vsql, ODBC, JDBC, and so on All cluster nodes with label `fusion:yes` and all remote database clients should be able to access this port.
5434	TCP	All database nodes	Used for intra-cluster and inter-cluster communication. Database opens the *database client port +1* (5434 by default) for intra-cluster communication. If the default *database client port +1* is not available, then database opens a random port for intra-cluster communication All database nodes should be able to access this port.
5438	TCP	All database nodes	Used for Management Console-to-node and node-to-node (agent) communication. For more information, see Changing Management Console or Agent Ports. All database nodes should be able to access this port.
5450	TCP	All database nodes	Used to connect to the Management Console from a web browser and allow communication from nodes to the Management Console application/web server. For more information, see Connecting to Management Console. All database nodes and web clients should be able to access this port.

Ports	Protocol	Node	Description
4803	TCP	All database nodes	Used for spread client connections. All database nodes should be able to access this port.
5433	UDP	All database nodes	Used for database spread monitoring. All database nodes should be able to access this port.
4804	UDP	All database nodes	Used for spread daemon to daemon connections. All database nodes should be able to access this port.
6543	UDP	All database nodes	Used to monitor spread daemon connections. All database nodes should be able to access this port.
4803	UDP	All database nodes	Used for spread daemon to daemon connection. All database nodes should be able to access this port.

CDF Master and Worker Node

CDF Management Portal

Port	Protocol	Node	Description
3000	TCP	Master	Used for accessing the CDF Management portal during CDF deployment from a web browser. This port is used to access the CDF Management portal only during CDF deployment. After deployment, port 5443 is used to access the CDF Management portal. Web clients must be able to access this port during the installation of CDF.
5443	TCP	Master	Used for accessing the CDF Management portal post CDF deployment from a web browser. Web clients must be able to access this port for administration and management of CDF.
5444	TCP	Master	Used for accessing the CDF Management portal post CDF deployment from a web browser, when using two-way (mutual) SSL authentication. Web clients must be able to access this port for administration and management of CDF, when using two-way (mutual) SSL authentication.

CDF

Ports	Protocol	Node	Description
8200	TCP	Master	Used by the `itom-vault` service which provides a secured configuration store. All cluster nodes should be able to access this port for the client connection.
8201	TCP	Master	Used by the `itom-vault` service which provides a secured configuration store. All cluster nodes should be able to access this port for peer member connections.

Kubernetes

Ports	Protocol	Node	Description
2380	TCP	Master	Used by the `etcd` component which provides a distributed configuration database. All the master nodes should be able to access this port for the `etcd` cluster communication.
4001	TCP	Master	Used by the `etcd` component which provides a distributed configuration database. All cluster nodes should be able to access this port for the client connection.
5000	TCP	Master	Used by `kube-registry` component which handles the management of container image delivery. All cluster nodes should be able to access this port to communicate with the local container registry.
8443	TCP	Master	This is a Kubernetes API server port. All cluster nodes should be able to access this port for internal communication.
8472	UDP	All nodes	Used by the Flannel service component which manages the internal cluster networking. All cluster nodes should be able to access this port for internal communication.
10250	TCP	All nodes	Used by the Kubelet service which functions as a local node agent that watches pod specifications through the Kubernetes API server. All cluster nodes should be able to access this port for internal communications and worker node Kubelet API for exec and logs.

Ports	Protocol	Node	Description
10251	TCP	All nodes	Used by `Kube-scheduler` component that watches for any new pod with no assigned node and assigns a node to the pod. All cluster nodes should be able to access this port for internal communication.
10252	TCP	All nodes	Used by `kube-controller-manager` component that runs controller processes which regulate the state of the cluster. All the cluster nodes should be able to access this port for internal communication.
10256	TCP	All nodes	Used by the `Kube-proxy` component, which is a network proxy that runs on each node, for exposing the services on each node. All the cluster nodes should be able to access this port for internal communication.

Transformation Hub

Ports	Protocol	Node with Label	Description
2181, 32181	TCP	zk:yes	Used by Kafka consumers like Kafka scheduler for database to connect to Zookeeper server. All database nodes must be able to access this port.
9093, 39093	TCP	kafka:yes	This is a SSL port used by Kafka broker to listen for incoming client connections. All data sources, such as SmartConnector, Identity Manager Driver for Entity Data Model, and Identity Governance should be able to access this port.
38080	TCP	th-platform: yes	Used by ArcSight Management Center (ArcMC) to connect to Transformation Hub. Also, used by the cluster nodes labeled `fusion: yes` to get the list of Kafka brokers. All cluster nodes with label `fusion: yes` and ArcSight Management Center should be able to access this port.

Identity Intelligence

Ports	Protocol	Node	Description
443	TCP	Master	Used for accessing the Identity Intelligence user interface from a web browser.

NFS

Ports	Protocol	Node	Description
111	TCP	NFS server	This is a NFS server port used by `portmapper` service. All cluster nodes should be able to access this port.
2049	TCP	NFS server	This is a NFS server port used by `nfsd` daemon. All the cluster nodes should be able to access this port.
20048	TCP	NFS server	This is a NFS server port used by `mountd` daemon All the cluster nodes should be able to access this port. **NOTE:** This port must be open only in a multi-node deployment.

SmartConnector

Ports	Protocol	Description
1515	TCP	Used by SmartConnector to receive events