Perceptive Content Server

Installation and Setup Guide

Version: Foundation EP1

Written by: Product Knowledge, R&D
Date: June 2022

Hyland™

Documentation Notice

Table of Contents

Install Perceptive Content

Running Perceptive Content on your network requires that you install Perceptive Content Server on a server computer and install at least one Perceptive Content Client on a computer that can access the server computer. Install the client on all computers on which a user performs Perceptive Content tasks, such as scanning and linking.

Perceptive Content Server supports 64-bit versions of Windows server 64-bit Linux Operating Systems. Perceptive Content can be set to use encrypted communication through TCP/IP to pass data between the server and clients. Each user accesses Perceptive Content from the client using a login ID and password. User authentication takes place on Perceptive Content Server, but you set it up in Perceptive Content Management Console.

An initial installation of Perceptive Content Client and Server requires you to complete several tasks in order. The following installation information assumes that you are performing an initial installation of Perceptive Content instead of upgrading from an earlier version of Perceptive Content. If you are updating Perceptive Content components, make sure you first refer to the *Update ReadMe* document. Sections of the update readme may reference procedures in this installation guide.

Perceptive Content components are not backwards compatible. For example, you must install version Perceptive Content Client Foundation EP1 (7.4.x) to work with version Perceptive Content Server Foundation EP1 (7.4.x). For product technical specifications and system requirements, refer to the *Technical Specifications* document.

To install the Perceptive Content environment, install the components in this order:

1. Install Erlang and RabbitMQ message queuing broker. Perceptive Content Server versions 7.1.4.x and higher depend on RabbitMQ, an open source software. See the RabbitMQ website for more information, and to download and install the product. For more information, see the **Appendix B: RabbitMQ considerations**.

    Important Prior to installing RabbitMQ and Erlang, review the Perceptive Content Server > Message Queuing specifications in the 7.1.x and higher *Perceptive Content Technical Specifications* for the supported versions of both products.

2. Install Perceptive Content Database. For more information, see the *Perceptive Content Database Installation and Setup Guide*.

3. Install Perceptive Content Data Source. For more information, see the *Perceptive Content Database Installation and Setup Guide*.

4. Install Perceptive Content Server.

5. Install Perceptive Content Client. For more information, see the *Perceptive Content Client Installation and Setup Guide*.

Important This document assumes you are installing Perceptive Content Server for the first time or that you have no earlier versions running on your computer. To update or upgrade from a previous version, see the *Perceptive Content Update Guide Foundation EP1* and *Perceptive Content Release Notes Foundation EP1*.

Using an active-active server environment

With an active-active environment, you can run multiple instances of Perceptive Content balanced across redundant nodes, also called clusters. One server environment is set up (the primary) on one node and at least one other server environment is created (secondary) on a different node. If the primary environment begins to fail, the system immediately switches over to the secondary environment. This offers two benefits. The first is that an active-active environment does not require the large-scale investment in extra hardware to backup data. The second benefit is that an active-active environment protects against system-wide crashes and avoids single points of failure because the server can switch to the secondary server environment if the primary server environment fails without having to shut down and restart the system.

You also can install the two parts of Perceptive Content Server, per-Server and shared Server, in independent locations. The per-Server is usually installed locally with respect to the installer. The shared Server piece, since it is only installed once per active-active environment, is usually installed on a network drive.

Note Whether you decide to install the per-server and shared server on the same machine or in different locations, they both must be installed.

Set up a server farm for load balancing

A server farm is a collection of real servers that operate behind a virtual IP address, streamlining server workload by spreading it among many physical servers using a load-balancer. For example, when a connection is made to a virtual IP address that is associated with a load-balancer, the load-balancer picks the best real server to handle the connection. A server farm also increases redundancy by allowing other servers to handle incoming requests if one fails.

A real server is a physical machine that hosts data, manages network resources, and processes workload from clients. Virtual servers are interfaces that accept incoming connections and route them to a real server. The system that the load-balancer uses to determine if a real server is available to accept incoming connections is called health monitoring.

Assembling a server farm is optional and is done after installing Perceptive Content Server and before installing Perceptive Content Client. For more information, see the steps in the Assemble and configure a server farm for Perceptive Content Server section in this document.

Note Setting up and configuring a server farm is optional. After installing Perceptive Content Server, continue with the steps in *Perceptive Content Database Installation and Setup Guide* and *Perceptive Content Client Installation and Setup Guide* if you are not setting up a server farm.

Installation Process

The following sections outline the high-level procedures that you need to perform to install and configure Perceptive Content and information you need to verify before the installation.

Prepare for the installation

Before you install Perceptive Content, verify the following information.

- Verify the installation checklist

- Obtain the TCP/IP host name or TCP/IP address and the authorization port (6000 is the default) of the computer on which you will install Perceptive Content Server. You need these to log on the first time.

- In Windows, verify that Microsoft TCP/IP is installed and configured. On Windows, make sure that you have an NTFS-formatted volume for all Perceptive Content Server executables and directory structures.

- On Linux, you need to know the root user password unless you plan to run Perceptive Content as a non-root user. For more information about running Perceptive Content as a non-root user, refer to the *Running Perceptive Content as a Non-Root User Best Practices Guide*.

- Make sure that you have sufficient disk space for executables and object storage.

- Decide whether this is a single server installation or if your environment requires an active-active failover installation.

- The server and client system time must be synced to GMT using third-party time server software.

Verify the installation checklist

This checklist describes the high-level procedures performed during this installation. Each of these procedures is described in detail later in this guide.

- Verify TCP/IP connectivity for Perceptive Content Server.

- Create users and groups on the operating system network domain or directory, if needed. User names and passwords in the Perceptive Content Management Console must match the network user names and passwords.

- Download and install the appropriate Perceptive Content Server for your operating system and database. We recommend that you create a new instance for the INOW database.

- Optional. Assemble a server farm.

- Start Perceptive Content Server.

- License Perceptive Content unless you are installing an evaluation copy.

- Start all Perceptive Content services.

- Create and test a login profile.

- Log in to Perceptive Content to create and test groups and users.

Verify TCP/IP Connectivity for Perceptive Content Server

At the Command Prompt window, type **ping <*server address*>** where <*server address*> is the IP address or the host name of the computer running Perceptive Content Server (for example, ping 206.18.19.25 or ping notesrvr) and then press ENTER.

If you are properly connected, the message you receive appears similar to the one below.

```
C:\>ping 206.18.19.25

Pinging 206.18.19.25 with 32 bytes of data:

Reply from 206.18.19.25: bytes=32 time=111ms TTL=240
Reply from 206.18.19.25: bytes=32 time=100ms TTL=240
Reply from 206.18.19.25: bytes=32 time=100ms TTL=240
Reply from 206.18.19.25: bytes=32 time=100ms TTL=240

Ping statistics for 206.18.19.25:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 100ms, Maximum =  111ms, Average =  102ms

C:\>
```

Note If you receive the message "Bad IP address" in response to the ping request, you are not connected to Perceptive Content Server. Check with your system administrator to verify the Perceptive Content Server IP address or name.

Open firewall and network ports for communication

To enable communications between Perceptive Content Server and Perceptive Content Clients and agents, you must open TCP ports on your network. Perceptive Content Server uses ports 6000, 5672, and 5671.

For best performance and stability, we recommend that you disable the firewall on the Perceptive Content Server computer. If you cannot disable the firewall, you must set exclusions in the firewall for ports 6000 for Perceptive Content Server processes.

- **Port 6000** - By default, port 6000 is used for communications between Perceptive Content Server and Perceptive Content Client, including communications with web servers hosting Perceptive Content web clients, such as Experience. If you need to change this port, modify the **inowd.port** setting in the \inserver\etc\inserver.ini file after installing Perceptive Content Server.

- **Port 5672** – Port 5672 is the commonly used port for RabbitMQ. Perceptive Content Server versions 7.1.4.x and higher depend on RabbitMQ for message queuing.

- **Port 5671** – TLS-encrypted communications with RabbitMQ default to port 5671. If TLS encryption is used for RabbitMQ, use that configured port instead of 5672.

IPv6 compatibility

Perceptive Content Server is capable of communicating with Internet Protocol, version 6 (IPv6) devices. IPv6 is the standard, global protocol by which IP addresses are issued to computers, servers, networks, and other Internet-compatible devices. The number of unassigned IP addresses for the previous protocol, IPv4, is diminishing. IPv6 was launched to provide a massive number of open IP addresses for future devices.

If your network has been set up to use IPv6 then you need to set the ipv6.enabled setting to TRUE. For more information, see the inow.ini file table.

Install Perceptive Content Server on Linux

Important This document assumes you are installing Perceptive Content Server for the first time or that you have no earlier versions running on your computer. To update or upgrade from a previous version, see the *Perceptive Content Update Guide*. Sections of the update guide may reference procedures in this installation guide. Depending on server size and speed, it can take more than two hours to complete this process. Before you proceed with the installation, perform the following actions.

- Verify that your version of Oracle matches the minimum specifications and patch levels listed in the *Technical Specifications* for this product.

- If you are running Linux, make sure that your environment includes C/C++ runtime libraries.

- In Linux, the account under which Perceptive Content Server process runs must have root privileges unless you are running Perceptive Content as a non-root user. For more information about running Perceptive Content as a non-root user, refer to the *Running Perceptive Content as a Non-Root User Best Practices Guide*.

- Verify that /etc/hosts lists the IP address, hostname, and any aliases for the Perceptive Content Server computer.

- You must log in as root or sudo to install Perceptive Content Server.

In the following steps, you create the main directory for Perceptive Content. The default directory name is inserver. It is important that you do not rename the inserver directory while it is running or after your configuration, and never install Perceptive Content in the root file system of the Linux server. The root file system does not have the capacity to accommodate Perceptive Content.

Before beginning the installation process, decide if you are going to install and configure Perceptive Content Server in a stand-alone server environment, or in an active-active server environment. A stand-alone server environment requires a second server for failover protection and optional data backups. An active-active server environment requires more than one server in order to run multiple instances of Perceptive Content Server for data redundancy.

Download the Perceptive Content Server files

To obtain Perceptive product installation files, contact the Hyland Software Technical Support group. For a list of Technical Support phone numbers, go to hyland.com/pswtscontact.

For a stand-alone server

In this procedure, you create an inserver directory, copy the tar file to that directory, and then extract the files to that directory.

1. To create the directory, navigate to the desired base location on the local machine, and enter **mkdir inserver**.

2. Copy **<*filename*>.tar.gz** into the directory you created in the previous step, replacing <*filename*> with the name of your Perceptive Content Server file.

3. To extract all the packaged files, navigate to the **inserver** directory and then type the following command and replace <*filename*> with the name of your Perceptive Content Server file.

   ```
   tar -xzvf <filename>.tar.gz
   ```

4. Execute the following command.

   ```
   chown -R root:bin *
   ```

5. Enter **ls -l** to verify that the directory contains the TAR file and the following directories: audit, bin, doc, envoy, etc, form, log, odbc, osm_01.00001, osm_02.00001, osm_03.00001, script, temp, update, user, workflow, and form. This directory also includes the environment.ini file.

   Note In the **inserver** directory, you can remove the **TAR** distribution file to save disk space.

For an active-active environment

In this procedure, you create an inserver directory, copy the tar file to that directory, and then extract the files to that directory.

1. To create the local **inserver** directory, navigate to the desired base location on the local machine, and enter **mkdir inserver**.

   Note The local directories on all of the machines in your environment should be in the same location due to the **odbc.ini** configuration settings.

2. To create the shared **inserver** directory, navigate to the desired base shared location, and enter **mkdir inserver**.

3. Copy the local **<ImageNow-Server-Active-Active-Local>.tar.gz** into the local **inserver** directory you created, replacing *<filename>* with the name of your Perceptive Content Server file.

4. Copy the shared **<ImageNow-Server-Active-Active-Shared>.tar.gz** into the shared **inserver** directory you created, replacing *<filename>* with the name of your Perceptive Content Server file.

5. Navigate to the local **inserver** directory and then type **tar –xzvf <ImageNow-Server-Active-Active-Local>.tar.gz** to extract all the packaged files for the local machine.

6. To extract all the packaged files for the shared machine, type **tar -xzvf <ImageNow-Server-Active-Active-Shared>.tar.gz**.

7. Open the **environment.ini** configuration file. Create this file if it does not exist.

8. In the **environment.ini** file, update the following settings.

   • [Directory Locations]

   • CONTENTDIR=[*location of the shared files*] this is the same shared inserver directory where you extracted the **<ImageNow-Server-Active-Active-Shared>.tar.gz** file.

9. Execute the following command.

   ```
   chown –R root:bin *
   ```

10. In the local **inserver** and shared **inserver** directories, enter **ls -l** to verify that the directory contains the following directories:

Location	Directories
local inserver	bin, log, odbc, temp, environment.ini
shared inserver	audit, doc, envoy, etc, form, osm_01.00001, osm_02.00001, osm_03.00001, script, workflow

11. Optional. In the local inserver and shared inserver directories, you can remove the appropriate archive file to save disk space:

Location	File

local inserver	ImageNow-Server-Active-Active-Local.tar.gz
shared inserver	ImageNow-Server-Active-Active-Shared.tar.gz

Update the inow.ini file

You must update the inow.ini file manually when installing Perceptive Content Server on Linux.

1. Navigate to the shared **/inserver/etc** directory and then open the **inow.ini** file with a text editor.

2. In the **odbc.dbms** setting, enter the appropriate database.

   Note The database type is case sensitive and should be entered as **SQLServer**, **Oracle**, or **PostgreSQL**. The default is **SQLServer**.

3. Change the **odbc.dsn** to the name of your Data Source (DSN), for example INOW. Ensure the section name is defined in the **odbc.ini** file.

4. Update **odbc.user.id** and **odbc.user.password** to your user id and password.

   Note The password supplied in the odbc.user.password setting is consumed by the application for encryption in the odbc.user.password.encrypted setting. This value is removed from the setting after encryption.

5. To update the message queuing server settings, complete the following substeps.

   1. Update **mq.host** to specify the hostname or IP address of the node running the message queuing broker.

   2. Update **mq.port** to specify the port the message queuing broker uses.

      Note The commonly used message queuing port for RabbitMQ is 5672.

   3. Optional. Update the **mq.vhost** with the name of the virtual host.

   4. Update **mq.username** and **mq.password** to the user name and password used when connecting to the message queuing broker.

      Note The password supplied in the mq.password setting is consumed for encryption in the mq.password.encrypted setting. This value is encrypted and removed from the mq.password setting after running the `inserver -encrypt-config` command. Do not manually update the mq.password.encrypted setting value.

   5. Optional. Update **mq.secure.enable** to TRUE to use SSL/TLS.

      Note For Linux, the SSL_CERT_DIR environment variable must be set to the location of the CAP.pem file.

6. To update the message queuing client settings, complete the following substeps.

   Note If any of the following settings are incorrectly specified in the inow.ini file, Perceptive Content Client cannot connect to the message queuing broker. If this occurs, a message is logged in Message Center.

   1. Optional. Update **mq.client.reconnect.interval** to specify, in seconds, how long Perceptive Content Client waits before reattempting to connect to the message queuing broker.

   2. Update **mq.client.host** to specify the hostname or IP address of the node running the message queuing broker.

3. Update **mq.client.username** and **mq.client.password** to the user name and password used when connecting to the message queuing broker.

   Note The password supplied in the mq.client.password setting is consumed for encryption in the mq.client.password.encrypted setting. This value is encrypted and removed from the mq.password setting after running the `inserver –encrypt-config` command. Do not manually update the mq.client.password.encrypted setting value.

4. Optional. Update **mq.client.secure.enable** to TRUE to use SSL/TLS.

7. Save and close the **inow.ini** file.

8. Source the **setenv.sh** file to set up your environment variables. For example, if you are using /bin/bash, the command to source the file is: **. ./setenv.sh**.

Create the Perceptive Manager user

In Perceptive Content, the Perceptive Manager user is the highest administrative user and is the top-level user in the system with access to change all security privileges. After you install Perceptive Content, you log on as the Perceptive Manager to set up additional users.

1. Navigate to the local **inserver/bin** directory.

2. Source the **setenv.sh** file. For example, if you are using /bin/bash, the command to source the file is: **. ./setenv.sh**.

3. Run the following command, substituting the name you want to use as the Perceptive Manager for *<username>*:

```
./intool --cmd create-bootstrap-user --username <username>
```

4. Execute the setup.sh file. For example, if you are using bin/bash, the command to source the file is:

```
../setup.sh
```

   Note This script initializes the OSM structure within the default install location, $CONTENTDIR. If you are using a different location for the OSM structure, modify all $CONTENTDIR references to the full path of your preferred location in the FSS section of the **setup.sh** file.

5. Choose file-system storage (FSS), which is your conventional directory structure. If you need to create EXT sets, you must run the necessary INTool commands. For more information, see the *Perceptive Content External OSM Plugin Best Practices Guide*.

Install Perceptive Content Server on Windows

Important This document assumes you are installing Perceptive Content Server for the first time or that you have no earlier versions running on your computer. To update or upgrade from a previous version, see the *Perceptive Content Update Guide*. Sections of the update guide may reference procedures in this installation guide.

Before you install, verify the following information.

- On the Perceptive Content Server computer, check the Windows Event Viewer to make sure that the computer has no DNS, hardware, or critical Windows errors.

- Verify that your system meets the requirements in the Product Technical Specifications. Then, verify your product compatibility outside of Perceptive Content, such as the compatibility between the service pack level of the operating system and your version of Microsoft SQL Server.

- The server and client system time must be synced to GMT using third-party time server software.

In addition, we recommend that you have the Microsoft SQL Server database running on a different computer than the computer with Perceptive Content Server.

Important When installed on a Windows Server operating system, all agents that read to or write from the OSMs must be set to run under a domain service account if you are using a remote storage device.

Download the Perceptive Content Server files

To obtain Perceptive product installation files, contact the Hyland Software Technical Support group. For a list of Technical Support phone numbers, go to hyland.com/pswtscontact.

Install Perceptive Content Server attended

Before beginning the installation process, decide if you are going to install and configure Perceptive Content Server in an active-passive server environment or in an active-active server environment. An active-passive server environment requires a second physical server for failover protection and data backups. An active-active server environment requires only one physical server to run multiple instances of Perceptive Content Server and for data redundancy.

Note Before proceeding, make sure that the INOW database is installed and online.

Install Perceptive Content Server for Windows

1. In the **Windows Explorer**, right-click the executable you downloaded and select **Run as Administrator**.

2. On the **Welcome to the Installation Wizard for Perceptive Content Server** page, click **Next**.

3. On the **License Agreement** page, review the information, scroll to the bottom of the agreement and click in the agreement field, click **I accept the terms in the license agreement** and then click **Next**.

4. On the **Setup Type** page, choose the following type of installation and then click **Next**.

    - **Stand-alone Instance** – for an active-passive server environment. The **Destination Folder** page appears.

    - **Active-Active Instance** – for an active-active server environment. The **Custom Setup** page appears.

5. If you selected a **Stand-alone Instance** and want to change the installation location from **C:\inserver**, in the **Destination Folder** page, click **Browse**. In the **Change Current Destination Folder** page, browse to the destination folder where you want to install the server files and then click **OK**. The **Destination Folder** page appears.

6. If you selected an **Active-Active Instance** and want to change the installation location, in the **Custom Setup** page, you must change both the **Per-Server Files** and the **Shared Server Files** instances using the following substeps.

    1. In the **Custom Setup** page, click **Per-Server Files** and then click **Change**.

    2. In the **Change Current Destination Folder** page, browse to the destination folder where you want to install the per-server files and then click **OK**. The **Custom Setup** page appears.

    3. In the **Custom Setup** page, click **Shared Server Files** and then click **Change**.

    4. In the **Change Current Destination Folder**, browse to the destination folder where you want to install the shared server files and then click **OK**.

7. Click **Next**.

8. On the **Perceptive Content Server Information** page, set your initial instance name and then click **Next**.

   Note For Active-Active installations, the **Initial instance name** text box is available for specifying the instance name. The instance name allows for multiple instances of the same agent or server to run in parallel in an active-active environment. Accept the default label or supply a different description for the initial instance of the service. You can enter a maximum of 40 characters. The following characters are not valid: \ / : * ? " < > |.

9. On the **Configure Perceptive Content Database** page, select your preferred database, specifiy your DSN, verify the DSN exists, enter the user name and password for the database, and then click **Next**.

10. On the **Select Storage Type** page, select the storage type and then click **Next**.

11. On the **Perceptive Content Setup** page, set the following options.

    - **Port number** – Specify the port number you want Perceptive Content to use to connect to the Perceptive Content Server, which is typically 6000.

    - **Perceptive Manager** – Accept the administrator as the default Perceptive Manager or supply a different user for this role.

    - Optional. **Language** – Select the language you want to use with Perceptive Content.

12. Click **Next**. When prompted, click **Yes** to confirm the default Perceptive Manager.

13. On the **Server-Side Configuration for RabbitMQ** page, configure the settings according to your RabbitMQ instance and click **Next**.

14. On the **Client-Side Configuration for RabbitMQ** page, configure the settings according to your RabbitMQ instance and click **Next**.

15. On the **Ready to Install the Program** page, click **Install**.

    The **ODBC Wire Protocol Driver Setup** page appears. Perform one of the following actions for your specific Driver.

    In the **ODBC Oracle Wire Protocol Driver Setup** dialog box, on the **General** tab, specify the following information.

    - **Description** – Optional. Enter a description of the DataSource.

    - **Host** – Enter the IP address for the database server.

    - **Port Number** – Enter the server port number.

    - **SID** – Enter the security identifier.

    In the **ODBC SQL Server Wire Protocol Driver Setup** dialog box, on the **General** tab, specify the following information.

    - **Description** – Optional. Enter a description of the DataSource.

    - **Host Name** – Enter the IP address for the database server.

    - **Port Number** – Enter the server port number.

    - **Database** – Enter the database name.

    - In the **ODBC PostgreSQL Wire Protocol Driver Setup** dialog box, on the **General** tab, specify the following information.

        - **Description** – Optional. Enter a description of the DataSource.

- • **Host Name** – Enter the IP address for the database server.

- • **Port Number** – Enter the server port number.

- • **Database Name** – Enter the database name.

- • On the Advanced tab, set the following information.

  - • **Extended Options** – Set the **StaticCursorLongColBuffLen** value to a number, in bytes, that is larger than your largest Learnmode script.

    Note This value should not be smaller than 8192. It is recommended that you set the initial value to 131072.

  - • Set the **Transaction Error Behavior** value to **2 – Rollback to Savepoint.**

  - • Set the **Report Codepage Conversion Errors** value to **1. – Return Error**.

  - • To connect the DSN via iScript, check the **ExtendedColumnMetadata** box.

16. Click **OK**.

17. On the **Installation Wizard Completed** page, perform the following substeps.

    1. Select the **Show the readme file** check box.

    2. Optional. If the **Show the Windows Installer log** check box appears, you can select the check box to view the log fie.

    3. Click **Finish**.

18. If you are prompted to restart, click **Yes**.

Install Perceptive Content Server unattended

Installing Perceptive Content Server silently is an automatic way to run unattended installations. If you follow the procedures below, you will not install Perceptive Content Server using a standard InstallShield interface. Using this silent installation method, you can do a custom installation or use a combination of default and customized settings.

Run the unattended installation

1. Set up your argument values to customize the unattended installation. If you do not manually set argument values, then the default values are used during the installation.

Argument	Description	Default	Example
L*v	This value is optional. If you use this argument, setup does not create directories. The path for the log file generation must be a valid, existing path. This argument is typically used to diagnose installation errors.	%TEMP%\\ImageNow Server*.log	/L*v C:\logs\client-install.txt

Argument	Description	Default	Example
INSTALLDIR	The default and recommended installation directory is [drive:]\inserver.	[drive:]\inserver	INSTALLDIR=C:\inserver
INSTALLDIR2	This is the install directory of your common files in an active-active environment.	The value of INSTALLDIR	INSTALLDIR2=S:\inserver
IND_INOWD.PORT	Perceptive Content Server port.	6000	IND_INOWD.PORT=7000
IMAGENOW_OWNER	The default Perceptive Manager ID.	administrator	IMAGENOW_OWNER=jdoe
INSTANCE_NAME	The name you give the initial instance of the service.	Primary	INSTANCE_NAME= "production inst"
IN_SUPPORTED. LOCALE	The language code that represents the language of Perceptive Content you are installing if you are using a localized version of Perceptive Content.	en	IN_SUPPORTED.LOCALE=de
OSM_STORAGE_TYPE	The OSM storage device type FSS (File System Storage).	FSS	OSM_STORAGE_TYPE=FSS
ODBC.DBMS	Sets the database management system for Perceptive Content. This value should be set to either SQLServer or Oracle.	SQLServer	ODBC.DBMS=Oracle
ODBC.DSN	The ODBC name.	Perceptive Content	ODBC.DSN=\"Perceptive Content\"
PORT_NO	The port number of the database.	1433	PORT_NO=1521

Argument	Description	Default	Example
ODBC_DBNAME	The name of the Perceptive Content SQL Server database. Only applicable for SQL Server installations.	INOW	ODBC_DBNAME=DBNAME
ODBC_SID	The name of the Perceptive Content Oracle database. Only applicable for SQL Server installations.	INOW	ODBC_SID=DBNAME
ODBC_DESCRIPTION	The ODBC description.	Perceptive Content	ODBC_DESCRIPTION=\"my description\"
IS_SQLSERVER_ SERVER	The hostname of the database server.	localhost	IS_SQLSERVER_SERVER=local host
MQ.HOST	The hostname of the RabbitMQ instance.	localhost	MQ.HOST=127.0.0.1
MQ.PORT	The port number of the RabbitMQ instance.	5672	MQ.PORT=5673
MQ.USERNAME	The RabbitMQ username.	guest	MQ.USERNAME=admin
MQ.PASSWORD	The RabbitMQ password.	guest	MQ.PASSWORD=1234
MQ.SECURE.ENABLE	This setting enables SSL on the RabbitMQ instance.	FALSE	MQ.SECURE.ENABLE=FALSE
MQ.VALIDATE.SERVER. CERTIFICATE.ENABLE	This setting validates the RabbitMQ SSL certificate.	TRUE	MQ.VALIDATE.SERVER.CERTI FICATE.ENABLE=TRUE
MQ.CLIENT.PORT	The port number the Client uses to connect to RabbitMQ.	5672 On upgrade, defaults to the previously configured MQ.PORT	MQ.CLIENT.PORT=5672

Argument	Description	Default	Example
MQ.CLIENT.SECURE.ENABLE	Enables the Client to connect to RabbitMQ using SSL.	FALSE	MQ.CLIENT.SECURE.ENABLE=FALSE
MQ.CLIENT.VALIDATE.SERVER.CERTIFICATE.ENABLE	If SSL is enabled for the Client, this setting enables validation of the SSL certificate.	TRUE	MQ.CLIENT.VALIDATE.SERVER.CERTIFICATE.ENABLE=TRUE

2. Enter the following command. You can use one of the commands in a **Command Prompt** window, in the provided batch file (Server_SilentInstall660.bat), or create a command line script for your deployment software.

```
ImageNowServer-ExternalDB.exe /s /V"<argument list>"
```

The following example shows the command with a defined argument list.

```
ImageNowServer-ExternalDB.exe /s /V"/qn /L*v \"C:\logs\server-install.txt\"
INSTALLDIR=\"C:\inserver\" ODBC.DBMS=SQLServer ODBC.DSN=\"my dsn name\"
ODBC_DESCRIPTION=\"my odbc description\" ODBC_DBNAME=INOW PORT_NO=1521
IND_INOWD.PORT=6000 IMAGENOW_OWNER=administrator IN_SUPPORTED.LOCALE=en
IS_SQLSERVER_SERVER=localhost INSTANCE_NAME=Production"
```

3. Verify your installation.

Increase performance for Perceptive Content

The following recommendations can increase performance after you have Perceptive Content running in production. In addition, as your user base gets larger, you can make the following changes to maximize the performance of your system.

In your anti-virus application, disable on-access scanning for the \inserver directory, including all subdirectories, and your database. When you use on-access scanning, your virus scanner continually examines Perceptive Content Server memory and file system, which can decrease performance. If you move any directory outside of \inserver, make sure you disable on-access scanning in the new location. You can verify the location of Perceptive Content directories in the \inserver\etc\inow.ini file.

Verify that the num.workers setting in the inserver.ini file is set to reflect an accurate number of users. The suggested ratio is to set one thread for every ten users.

Migrate the shared directory

Perceptive Content Server installer creates multiple entries in the registry that point to the location of the shared configuration for the active-active environment. If this location was modified in the environment.ini file, it is necessary to move the shared directory to the new location. Complete the following steps as an administrator to move the shared directory.

1. Back up the shared directory.

2. On each server, complete the following substeps.

1. Open **regedit.exe**, navigate to HKEY_LOCAL_MACHINE\SOFTWARE\Perceptive Software\ImageNow\Server, and delete the `SharedInstallPath` string value.

2. Copy the `SharedDirectoryMigration.bat` file to the directory where the INServer installer is located.

   Note The installer must correspond to the installed version of INServer.

3. Open a **Command Prompt** window and navigate to the directory where your installer is located.

4. Run the batch file using the following command.

   ```
   SharedDirectoryMigration.bat [serverinstaller.exe] [newshareddirectorylocation]
   ```

3. Move the shared directory to the new location.

Configure trust stores for SSL/TLS on Windows

Perceptive Content Server's environment must be configured to trust CAs used by several different functionalities. Windows agents use the Windows Certificate Manager, while Linux agents must be configured to use an external trust store.

To import a certificate into Perceptive Content Server on a Windows operating system, complete the following steps.

1. Click **Start** > **Run**.

2. In the **Run** dialog box, type **MMC**, and then click **OK**.

3. In the left pane, in the **Console Root** tree, click Certificates (Local Computer) > Trusted Root Certification Authorities.

4. Right-click anywhere in the right pane, and click **All Tasks** > **Import**.

5. In the **Certificate Import Wizard** welcome page, click **Next**.

6. In the **File to Import** page, browse to the certificate you just created, and click **Next**.

7. In the **Certificate Store** page, verify **Place all certificates in the following store** is selected, and click **Next**.

8. In the **Completing the Certificate Import Wizard** page, click **Finish**.

9. Double-click the certificate and verify it imported successfully.

Configure trust stores for SSL/TLS on Linux

Since the location of the trust store can be anywhere on Linux, environment variables may be configured to point to the correct location.

SSL_CERT_FILE

SSL_CERT_FILE should specify the full path to a file of CA certificates in PEM format. The file may contain any number of CA certificates in sequence with text allowed before, between, and after certificates. When this environment variable is not specified, it defaults to the file **cert.pem** in the default **OpenSSL** directory.

The following is an example of the file format.

```
# Description of certificate
-----BEGIN CERTIFICATE-----
... (CA certificate in base64 encoding) ...
```

```
-----END CERTIFICATE-----

# Description of certificate
-----BEGIN CERTIFICATE-----
... (CA certificate in base64 encoding) ...
-----END CERTIFICATE-----

# Description of certificate
-----BEGIN CERTIFICATE-----
... (CA certificate in base64 encoding) ...
-----END CERTIFICATE-----
```

SSL_CERT_DIR

SSL_CERT_DIR should specify the full path to a directory of CA certificates in PEM format. Each file should contain one CA certificate and have a unique CA subject name hash value. When this environment variable is not specified, it defaults to the **certs** directory in the default **OpenSSL** directory.

Note that the system searches **SSL_CERT_FILE** for a given certificate before searching **SSL_CERT_DIR**.

Assemble and configure a server farm for Perceptive Content Server

The following instructions provide a high-level overview for setting up and configuring a server farm using a Cisco Application Control Engine (ACE) Module. For more information and configuration parameters, see the appropriate Cisco documentation. If you are using a different third-party product to set up a server farm, refer to that product's documentation for similar instructions.

Assembling a server farm is comprised of the following steps.

- Setting up server health monitoring

- Setting up real servers

- Configuring the server farm

- Configuring a virtual server

Set up server health monitoring

The following instructions include how to set up two different probe methods. The first method uses specific probes that verify that the servers are running and responding to requests. The second method is a simple machine ping that pings the machine to verify that it is online.

Method one is the preferred method and you can implement both types of health checks in the same environment. However, you do not need to use both methods to set up server health monitoring.

Set specific health probes

Perform the following steps to set a specific health probe for Perceptive Content Server.

1. Open the **Cisco ACE** configuration tool.

2. Click the **Config tab**.

3. In the **Load Balancing** menu, select **Health Monitoring**.

4. In the **Probe name** field, provide a name for the probe.

5. In the **Health probe type** field, select **TCP** for the type.

6. In the **Port** field, enter `6000`. This is the port that Perceptive Content Server listens on, as defined in the **inserver.ini** file.

7. In the **Send Data** field, enter the following information.

   `000000011SERVER_PING000000008END_MARK000000013CLOSE_SESSION`

8. In the **Expect Regular Expression** field, under the **More Settings** header, enter the following information.

   `.*SUCCESS.*`

   Note The health probe ensures that `SUCCESS` appears in the returned string. A `SUCCESS` response indicates the server is connected to the database and accepting logins. If `.*` tokens are not valid on your load balancer, adjust the configured receive string. Load balancers that perform a 'contains' comparison with the receive string do not require the `.*` token.

9. In the **Probe interval count** field, enter a probe interval count. This is the time interval between sending probes during a health check. The recommended value is 15 seconds.

10. In the **Pass detect (failed probe) interval** field, enter a pass detect count. This is the time interval between sending probes during a health check when the server is in a known bad state. The recommended value is 30 seconds.

11. In the **Fail detect count** field, enter a fail-detail count. This is the consecutive number of times a probe must fail before the server is marked as failed. The recommended count is two times. If you set the Probe interval count to 15 seconds, the load balancer stops sending new connections to the server within 30 seconds of it going down.

12. In the **Pass detect count** field, enter a pass detect count. This is the number of successful responses a probe must produce before the server is marked healthy. The recommended count is two times. If you set the **Pass detect (failed probe) interval** to 30 seconds, the load balancer starts to send new connections to the server within one minute of it coming back up.

13. In the **Receive timeout for a response count** field, enter a receive timeout for a response count. This is the amount of time that a server has to return a response during a probe. If it does not return a response within the set time, the probe fails. The recommended value is five seconds.

Set a simple health verification

The following steps establish a simple machine ping that verifies the machine is online.

1. Open the **Cisco ACE configuration** tool.

2. Click the **Config** tab.

3. In the **Load Balancing** menu, select **Health Monitoring**.

4. Create a new health probe by clicking the **+** (plus) symbol.

5. In the **Probe name** field, provide a name for the probe.

6. In the **Health probe type** field, select **ICMP** for the type.

7. In the **Pass detect (failed probe) interval** field, enter a pass detect count. This is the time interval between sending probes during a health check when the server is in a known bad state. The recommended value is 30 seconds.

8. In the **Fail detect count** field, enter a fail-detect count. This is the consecutive number of times a probe must fail before the server is marked as failed. The recommended count is two times. For

example, if you set the Pass detect (failed probe) interval count to 15 seconds, the load balancer stops sending new connections to the server within 30 seconds of it going down.

9. In the **Pass detect count** field, enter a pass detect count. This is the number of successful responses a probe must produce before the server is marked healthy. The recommended count is two times. For example, if you set the Pass detect (failed probe) interval to 30 seconds, the load balancer starts to send new connections to the server within one minute of it coming back up.

10. In the **Receive timeout for a response count** field, enter a receive timeout for a response count. This is the amount of time that a server has to return a response during a probe. If it does not return a response within the set time, the probe fails. The recommended value is five seconds.

Set up real servers

Set up the real servers on a VLAN interface. Make sure that these servers exist on the same subnet as the VLAN interface. After setting up the servers on real machines, use the following steps to add them to the hardware load-balancer configuration.

1. Open the **Cisco ACE configuration** tool.

2. Click the **Config** tab.

3. Add a new server by clicking the **+** (plus) symbol.

4. In the **Server Name** field, enter the server name.

5. In the **IP address** field, enter the IP address.

6. In the **Connection limits** field, enter the connection limits.

Configure the server farm

Configure the server farm for Perceptive Content Server using the steps in the following sections.

Configure the server farm for Perceptive Content Server

1. Open the **Cisco ACE** configuration tool.

2. Click the **Config** tab.

3. In the **Load Balancing** menu, select **Server Farms**.

4. Add a new server farm by clicking the + (plus) symbol.

5. To monitor the health of the server farm, set the probe to the one that you configured when you set up server-health monitoring.

6. Select **Purge for the fail action**. This option sends a reset to terminate the socket connection when a server fails.

7. Click **Deploy Now**.

8. Under the **Real Server @ <ServerFarmName>** heading, add the real servers to the server farm by clicking the **+** (plus) symbol.

9. Complete the following substeps.

    1. In the **Real server name** list, select the real server name.

    2. In the **Port** field, set the port number to the port Perceptive Content Server listens on. This is the same port that was set when you created the health probe for Perceptive Content Server. The default is **6000**, as defined in the **inserver.ini** file.

3. Optional. Set a backup server and port. The backup server becomes active if the real server is in a failed state.

4. In the **Server weight** area, configure the server weight. Servers with higher weights receive more connections as a ratio of their weight to the other servers' weights.

5. Click **Deploy Now**.

10. Click the **Predictor** tab and select **Least Connections** for the predictor type.

Configure a virtual server for Perceptive Content Server

Associate the server farm with a virtual server using the following steps for Perceptive Content Server.

1. Click the **Config** tab.

2. In the **Load Balancing** menu, select **Virtual Servers**.

3. Add a new virtual server by clicking the **+** (plus) symbol.

   Note This requires administrative permissions on the Cisco ACE device.

4. In the **Provide a virtual server IP address** field, enter a virtual server IP address. This is the address clients use to connect to one of the machines in the server farm.

5. In the **Virtual Server** list, select **VLAN interface** for the virtual server.

6. In the **Load Balance** list, select **Load Balance** for the primary action.

7. Use the server farm you configured in the "Configure the server farm" section as the server farm.

Verify the server farm setup

Complete the following steps to ensure that you have successfully set up and configured your server farm.

1. Start **Perceptive Content Server** on the real server machines.

2. In the **Cisco Application Control Engine (ACE) Module**, select the **Monitors** tab.

3. In the left column, click **Real Servers**. Ensure that the servers you configured appear as In Service.

4. Continue installing Perceptive Content Client. For more information, reference the *Perceptive Content Client Installation and Setup Guide*.

Install a Perceptive Content Server license

Before entering your license, you must install the Perceptive Content Server and at least one Perceptive Content Client. For more information on installing Perceptive Content Client, refer to the *Perceptive Content Client Installation and Setup Guide.* Only a Perceptive Manager user can install the license. In addition, on Linux, to obtain the hardware information for the Perceptive Content Server, you must also be the root user.

Obtain the license files

To obtain the system fingerprint for Perceptive Content Server, complete the following steps. Only a Perceptive Manager user can complete this task.

1. Start **Perceptive Content Server**. If you have an active-active setup, you must start all instances of Perceptive Content Server on all nodes.

2. Generate a system fingerprint using the following substeps.

   1. Click **Start** and select **All Programs** > **Perceptive Content** > **Perceptive Content Management Console**.

   2. In the **Login** page, click **License Manager**.

   3. In the **License Management** dialog box, select **Save system fingerprint** and click **OK**.

   4. In the **Save As** dialog box, enter a name for the file and then navigate to the location where you want to save the report. Click **Save**.

3. Contact your Perceptive Software representative for instructions on where to send the system fingerprint file to obtain your license. The system fingerprint file has a SYSFP extension.

4. When you receive the license files, store the license files in a temporary directory on the Perceptive Content Server computer.

Install Perceptive Content product licenses

Before entering your licenses, you must have installed the Perceptive Content Server and at least one Perceptive Content Client. Only a Perceptive Manager user can install the license. The Perceptive Content Client must be available on a Windows machine in order to install the Perceptive Content product licenses.

1. When you receive the license files from your Perceptive Software representative, copy them to a temporary folder where you can access them with a Perceptive Content Client.

2. Upload licenses, as explained in the following substeps.

   1. Click **Start**, point to **All Programs**, and then select **Perceptive Content**.

   2. In the login page, click **License Manager**.

   3. In the **License Management** dialog box, select **Upload Licenses** and click **OK**.

   4. Navigate to the folder where you stored the Perceptive Content license files, select the LIC files to upload, and click **Open**.

   5. Enter the **User Name**, **Password**, and **Server Location** and click **OK**.

   6. Optional. The **License Upload** dialog box lets you view the type name, actual license code, and current status of each license upload. To display detailed information for a specific license, select the appropriate row.

3. Click **OK**.

Obtain the license files for Amazon EC2

To obtain the license files required to run Perceptive Content on an Amazon EC2 compute resource, complete the following steps.

1. Stop all Perceptive Content services.

2. In the **/<path>/inserver/etc** directory, backup any **in_hwfp.<node>** files to a different directory.

   Note You can delete this file after you confirm that the installation procedure was successful.

3. Open the **inow.ini** file with a text editor and locate the **Licenses** section.

4. Set **hardware.amazonec2.support** to TRUE.

5. Save the **inow.ini** file and close the text editor.

6. Complete the steps in the Obtain the license files and Install Perceptive Content product licenses sections.

Obtain additional license files

If you have a licensed Perceptive Content environment and are obtaining additional licenses, such as for an active-active server environment, complete the following steps to generate a system fingerprint.

1. Start all Perceptive Content services.

2. Generate a new system fingerprint using the following substeps.

   1. Click **Start** and select **All Programs** > **Perceptive Content**.

   2. In the **Login** page, click **License Manager**.

   3. In the **License Management** dialog box, select **Save** system fingerprint and click **OK**.

   4. In the **Save As** dialog box, enter a name for the file and then navigate to the location where you want to save the report. Click **Save**.

3. Contact your Perceptive Software representative for instructions on where to send the system fingerprint file to obtain your license. The system fingerprint file has a SYSFP extension.

4. When you receive the license files, store the license files in a temporary directory on the Perceptive Content Server computer, and then install the licenses.

Disable anti-virus scanning of the Perceptive Content directory

In your anti-virus application, disable on-access scanning **for the \inserver** directory, including all subdirectories.

Note When you use on-access scanning, your virus scanner continually examines the Perceptive Content Server memory and file system, which can decrease performance. If you move any directory outside of \inserver, make sure you disable on-access scanning in the new location. You can verify the location of Perceptive Content directories through the \inserver\etc\inow.ini file.

Start all Perceptive Content services

After everything is licensed, you can start the Perceptive Content services. Skip this procedure if your services are running.

Start services using a terminal

1. If this is a new Linux terminal session, source the **setenv.sh** file to set up your environment variables.

   For example, if you are using /bin/bash, the command to source the file is: **. ./setenv.sh**.

   Note If you ran **setenv.sh** for this session, **do not** run it again. If you are unsure whether you have run the command, use the following command to check the CONTENTDIR environment variable.

   ```
   echo $CONTENTDIR
   ```

2. Start all services and agents.

   For example, to start the main services and agents (inserver, inserverAlarm, inserverBatch, inserverEM, inserverFS, inserverImp, inserverJob, inserverMonitor, inserverNotification, inserverOSM, inserverTask and inserverWorkflow), navigate to the **inserver/bin**, or **inserver/bin64** directory (depending on the bitness of the product), and then execute the following commands.

```
inserver -start [instance_name]
inserverAlarm -start [instance_name]
inserverBatch -start [instance_name]
inserverEM -start [instance_name]
inserverFS -start [instance_name]
inserverImp -start [instance_name]
inserverJob -start [instance_name]
inserverMonitor -start
inserverNotification -start [instance_name]
inserverOSM -start [instance_name]
inserverTask -start [instance_name]
inserverWorkflow -start [instance_name]
```

Notes

- The instance name allows for multiple instances of the same agent or server to run in parallel in an active-active environment. If you do not specify an instance name, the default value is set to **Primary**.

- If you purchased additional Perceptive Content agents or services you must also start them. Each time you log on, you must set your environment variables, as shown in the previous step, before you execute startup commands like `` `pwd`/inserver``.

Using Windows Computer Management

1. On your Windows Desktop, right-click the **My Computer** shortcut and select **Manage**.

2. In the **Computer Management** dialog box, click **Services and Applications**.

3. Click **Services**.

4. In the right pane, right-click the service you want to start and select **Start**.

Appendix A: IMAGENOWLOCALDIR6 and IMAGENOWDIR6 environment variables

The IMAGENOWLOCALDIR6 environment variable is a variable that holds the absolute file path to ImageNow Server resources in the inserver directory on a local machine. This is required for any ImageNow Server environment. The IMAGENOWDIR6 variable holds the absolute file path to the ImageNow Server resources in the inserver directory on a shared device. It is only required if running an active-active server environment.

IMAGENOWLOCALDIR6 environment variable

This environment variable is required and specifies the location of the following folders. These folders cannot be shared among the different instances of ImageNow Server.

- bin

- bin64

- temp

- log (default)

- ODBC

Note The log directory is included in this location by default, but it may be moved to the shared location.

IMAGENOWDIR6 environment variable

The IMAGENOWDIR6 environment variable is only required in an active-active environment. Some of the ImageNow Server resources must be moved to a shared file system, so they can be accessed by all of the different machines in an active-active environment. On Windows, this needs to be a NTFS file share accessible by UNC path. On Linux, the shared file system needs to be mounted locally and accessed like any other local directory. The path to that share is saved in the IMAGENOWDIR6 variable. If this environment variable is not present, its value is defaulted to the value of IMAGENOWLOCALDIR6. The environment variable is pulled on server startup from the environment.ini file located in the inserver directory on the local machine.

IMAGENOWDIR6 should specify the location of the following folders:

- envoy

- etc

- form

- job

- learnmode

- osm

- script

Appendix B: RabbitMQ considerations

Perceptive Content Server and all Perceptive agents, versions 7.1.4.x and higher, depend on RabbitMQ, a third-party message queuing broker. RabbitMQ provides industry-leading message brokering for Perceptive for high throughput and guaranteed delivery of messages. RabbitMQ is an open source software that requires Erlang, a programming language. You must download and install the supported versions of both products prior to upgrading to Perceptive Content 7.1.4.x. For further considerations, see the following list.

- Perceptive Content relies on the AMQP 0.9.1 protocol from RabbitMQ, which is enabled by default.

- When utilizing Active-Active mode for Perceptive Content, you must configure multiple cluster nodes of RabbitMQ to achieve high availability for the system. We highly recommend that you follow RabbitMQ guidelines and configure clusters with an odd number of nodes. For more information, see the RabbitMQ Clustering Guide.

- You are not required to share the same operating system or host machine for Perceptive Content and RabbitMQ. We recommend separate resources for RabbitMQ cluster nodes for the best high-availability and workload balancing.

- You must configure RabbitMQ for privacy over data connections. For more information, see the RabbitMQ Configuration Guide.

Note For more information on RabbitMQ, visit the RabbitMQ documentation site.

Appendix C: JSON Attribute Paths

The **sso.openid.profile.<profileName>.user.claim** setting for OpenID Connect configuration uses a JSON attribute path. This is the path to the claim that contains username values within a JSON payload. This is a period delimited grammar, that can be used to traverse into nested JSON objects.

Simple mapping

```
{
    "username" : "hylanduser",
    "sub" : "263DFEBC-EB43-46E7-A4D7-398C5D161190"
}
```

In most cases, the mapping should be simple. In the above example, the user.claim setting would be **username**.

Escaping characters

```
{
    "user.name" : "hylanduser",
    "sub" : "263DFEBC-EB43-46E7-A4D7-398C5D161190"
}
```

If you have a claim with a period . or square bracket ([or]) in the name. Then a backslash (\) must be used to escape the character. The user.claim for the set of claims above would be **user\.name**.

Nested object mapping

```
{
    "sub" : "263DFEBC-EB43-46E7-A4D7-398C5D161190",
    "user" : {
        "name" : "hylanduser"
    }
}
```

For nested JSON objects, inner objects are accessed through a period. In the above example, the user.claim setting would be **user.name**.

Array element mapping

```
{
    "sub" : "263DFEBC-EB43-46E7-A4D7-398C5D161190",
    "user" : ["hylanduser", "hylanduser@onbase.net"]
}
```

For JSON array elements, inner values are accessed through square brackets. In the above example, the user.claim setting would be **user[0]**.

```
{
    "sub" : "263DFEBC-EB43-46E7-A4D7-398C5D161190",
    "user" : [
        {
            "type" : "name",
            "value" : "hylanduser"
        },
        {
            "type" : "email",
            "value" : "hylanduser@onbase.net"
        }
    ]
}
```

For JSON arrays of objects, additional logic may be used for equality using the operator eq inside the square brackets. The left hand side of the eq operator must be a JSON property name, and the right hand side must be a value to compare against the property's value. The eq operator supports numeric, boolean, and string equality. JSON property names must be escaped if they contain whitespace (" "), quotation marks ("), or apostrophes ('). JSON string values must be escaped if they contain quotation marks ("). In the above example, the **sso.openid.profile.<profileName>.user.claim** setting would be **user[type eq "name"].value**.