

Perceptive Content Database

Installation and Setup Guide

Version: 7.2.x

Written by: Product Knowledge, R&D

Date: August 2019

Table of Contents

Perceptive Content Database prerequisites.....	3
Download the Perceptive Content Database files.....	3
Verify Microsoft SQL Server instance prerequisites.....	3
Create a Microsoft SQL Server INOW database	3
<i>Connect to your SQL Server.....</i>	<i>3</i>
<i>Determine the database collation for a Unicode environment.....</i>	<i>4</i>
<i>Create the INOW database</i>	<i>4</i>
<i>Set up SQL Server to use Secure Sockets Layer (SSL) security.....</i>	<i>5</i>
Create an Oracle INOW database	6
<i>Additional database components.....</i>	<i>6</i>
<i>Default Tablespaces</i>	<i>8</i>
<i>About the INUSER Database Schema (User).....</i>	<i>9</i>
<i>Create the database objects.....</i>	<i>9</i>
<i>Set up Oracle to use Secure Sockets Layer (SSL) security.....</i>	<i>10</i>
Create a PostgreSQL INOW database	10
<i>Locating the PostgreSQL configuration files.....</i>	<i>10</i>
<i>Configure PostgreSQL to allow remote connections.....</i>	<i>10</i>
<i>Configure connection settings</i>	<i>11</i>
<i>Default Tablespaces</i>	<i>11</i>
<i>Create the INOW database</i>	<i>12</i>
Database access through INEMUSER.....	12
<i>Order processing of external messages.....</i>	<i>12</i>
<i>Unlock the INEMUSER account.....</i>	<i>13</i>

Perceptive Content Database prerequisites

This document assumes that you are installing the Perceptive Content Database for the first time and that you are not already running an earlier version of Perceptive Content.

Before you install, verify that your system meets the requirements in the [Product Content 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 the version for your specific relational database management system (RDBMS). Supported RDBMS platforms include SQL Server, Oracle and PostgreSQL.

Prior to installation, we recommend reading the *Perceptive Content Best Practices Guide* for the system you are installing.

Download the Perceptive Content Database files

To obtain Perceptive product installation files, go to [Perceptive Software Downloads](#) or contact the Hyland Software Technical Support group. For a list of Technical Support phone numbers, go to hyland.com/pswtscontact.

The following database creation scripts are included in the Perceptive Content Server download files.

- SunflowerSS.sql (ANSI)
- SunflowerSS_Unicode.sql (Unicode)
- SunflowerORA.sql (ANSI Only)
- SunflowerPG.sql (Unicode)

Verify Microsoft SQL Server instance prerequisites

1. In **SQL Server Management Studio**, in the **Object Explorer** pane, right-click your server and then click **Properties**. Typically, your server is the first item listed in the tree.
2. In the **Server Properties** dialog box, under **Select a Page**, click **Security**.
3. Under **Server authentication**, select **SQL Server and Windows Authentication Mode**, and click **OK**.
4. If prompted, click **Yes** to restart the server.

Important Restarting the server shuts down and starts up all databases on this instance.

5. In **SQL Server Configuration Manager**, ensure the TCP/IP Protocol is enabled for the IP address and port you intend to use for your configuration.

For more information on configuring your instance of SQL Server and INOW database for optimal performance, refer to the [Microsoft SQL Server Best Practices Guide](#).

For more information on configuring encrypted communication using SSL between the database and Perceptive Content Server, refer to the [Perceptive Content Server Installation and Setup Guide](#).

Create a Microsoft SQL Server INOW database

Follow the steps in the following sections in order to create the INOW database.

Connect to your SQL Server

1. Log in to the **SQL Server computer** with a user ID that has local administrator rights, click the **Start**

button, point to **Programs**, point to your Microsoft SQL Server version, and then click **SQL Server Management Studio**.

2. In the **Connect to Server** dialog box, enter the information in the **Server Type**, **Server Name**, and **Authentication** fields. Click **Connect**.
3. Continue with the following steps as needed for your environment.

Determine the database collation for a Unicode environment

If you are installing the Unicode version of Perceptive Content then you must manually specify the appropriate default collation for your database. Be mindful of the tempdb database collation when selecting the default collation as incompatibilities between collations, in some cases, can cause unexpected errors and require the use of expression-level collations.

For Unicode installations, we use the nvarchar data type for character data. The collation you choose further defines the character encoding to either UTF-16 or UCS-2 encoding. The collation also determines the rules applied to character data with regard to evaluating characters within SQL expressions such as string compares for data access and filtering and sorting of results.

For Unicode installations, we recommend that you select a collation that best fits the primary language you are using to drive the business. If you are primarily using English to conduct day-to-day operations then you may want to use the Windows collation Latin1_General_100_CI_AS (code page 1252) for the default database collation.

Important

- We recommend you avoid using Binary collations (*_BIN*) because they do not support CI (case insensitive searching) and could lead to undesirable results within Perceptive Content.
- Perceptive Content is designed to work with case-insensitive collations. Choose a collation that includes the _CI option.

To create the Perceptive Content database for a Unicode environment complete the following steps to set the SQL Server database default collation.

1. Open the Perceptive Content Database SunflowerSS_Unicode.sql creation script.
2. In **SQL Server Management Studio**, display a list of all of the supported collations on your system by executing the following SQL query.

```
SELECT * FROM fn_helpcollations() WHERE NAME LIKE '%_CI%' ORDER BY NAME;
```

3. At the beginning of the DDL script, locate the following collation property:

```
COLLATE REPLACE_WITH_VALID_COLLATION_NAME
```

4. Change the collation property to the appropriate language from the list that you displayed in step 2. For example, to set the collation to be generic for western European languages, change the property to the following:

```
Latin1_General_100_CI_AS
```

Create the INOW database

Use the following steps to create the INOW database on Microsoft SQL Server.

Important Verify the data directory for Microsoft SQL before creating the database. This is the directory where the database files appear after you execute the following SQL file. If you do not want the location of the INOW database files to reside in the default data directory, use the appropriate locations for each of the three file names.

Depending on the environment type you are creating, SunflowerSS.sql in the steps below can refer to SunflowerSS.sql or SunflowerSS_Unicode.sql.

1. In **SQL Server Management Studio**, in the **Object Explorer**, select the **master** database. The master database is located under the **System Databases** folder.
2. In the **File** menu, select **Open** and then select **File**
3. In the **Open File** dialog box, navigate to the location that contains the **SunflowerSS.sql** file, select the SQL file, and then click **Open**.
4. If the system prompts you to connect to the database, enter credentials for a user that has the sysadmin server role, and select **Connect** to make the connection.
5. In the **SunflowerSS.sql** file, update the paths specified in the **Create Database** and the **Alter Database** commands to indicate the desired location of the database files.

You must change the path after 'FILENAME = N' to the full path (drive, directory, and filename with extension) for all three files.

Notes

- The directories you specify must already exist.
 - The Primary file group (INOW_Data.mdf) contains all the tables (clustered indexes).
 - The Secondary file group (INOW_Index.ndf) contains all the non-clustered indexes.
 - The Log file (INOW_Log.ldf) contains records of modifications made by each transaction.
6. To execute the query batch file, in the **Query** menu, select **Execute**. The execution is finished when the completion message appears at the bottom of the **Query** window.

Set up SQL Server to use Secure Sockets Layer (SSL) security

SSL security uses cryptography and symmetric encryption to provide communication security at the transport layer for data sent over a network for application-specific protocols such as HTML and SMTP.

To use SSL security, you must import a certificate from a third party Certificate Authority, bind the certificate to a port number, and configure the certificate. Use the following steps to set up SSL security.

1. Click **Start**, and then click **Run**.
2. In the **Run** dialog box, type **cmd** and then click **OK**.
 1. In the **Command Prompt** window, run **mmc.exe**.
 2. To add the certificates snap-in for the local machine, click **File > Add/Remove snap-in > Certificates**.
 3. Select **Computer account**, and then click **Finish** and **OK**.
3. In the tree view, navigate to **Certificates > Personal** and right-click **Certificates**.
 1. Select **All Tasks > Import**
 2. In the **Certificate Import Wizard** dialog, browse to the **PFX** file and click **Next**.
 3. To import the key, enter the password, select the options you want, and click **Next**.

To secure your PFX file somewhere other than on the production machines where it is used, do not select **Mark this key as exportable**.

4. When prompted, specify where the certificates are stored, and select **Personal**.

5. Click **Finish**.
4. Navigate to the file location of the imported certificate. In the **Actions** pane, under the certificate name, click **More actions**, then **All tasks**, and then **Manage private keys**.
5. Add the instance name that the SQL Server instance is running as, and then give the user full control.
6. Open **ODBC Datasource Administrator** and complete the following sub-steps.
 1. On the **DSN** tab, select **Datasource**, and then click **Configure**.
 2. On the **Security** tab, change the encryption method from **none** to **SSL**.
 3. To verify the connection, click **Test Connection**.
7. Close **ODBC Datasource Administrator**.
 - For **Windows Vista and higher**, use the **Netsh.exe** tool, as shown in the following example.

```
netsh http add sslcert ipport=0.0.0.0:8000
certhash=0000000000003ed9cd0c315bbb6dc1c08da5e6 appid={00112233-4455-6677-8899-
AABCCDDEEFF}
```

The certhash parameter specifies the thumbprint of the certificate, and the ipport parameter specifies the IP address and port and functions just like the -i switch of the Httpcfg.exe tool described. The appid parameter is a GUID that identifies the owning application.

8. Restart the **SQL Server instance**.

Create an Oracle INOW database

To create an Oracle database, we recommend that you adhere to your organizations best practices with regard to storage and database file placement. The Perceptive Content database supports the use of Exadata or pluggable databases (PDB). If your organization does not have best practices defined for Oracle databases then please refer to the [Using Perceptive Content with an Oracle Database Best Practices Guide](#).

Additional database components

No additional database components are required for the Perceptive Content Database; however, if you intend to use SSL or Oracle Net native encryption for data encryption and integrity you must purchase and install the Advanced Security option. Refer to the [SSL Security with Oracle Database Best Practices Guide](#) for assistance with configuring two-way SSL authentication and communication between the Perceptive Content Database and Application Server.

Create the INOW database

Depending on the version of Oracle you are using, there may be additional configuration options not found in this guide. Please refer to the best practices for your organization.

1. Using the **Oracle Database Configuration Assistant (DBCA)**, create a new database by selecting the **Custom Database** template, click **Next** and name it according to the naming conventions for your organization. For example, **INOW**.
2. On the **Management Options** page, follow the best practices set by your company for configuring Enterprise Manager and Automatic Maintenance Tasks.
3. On the **Database Credentials** page, follow the best practices set by your company.
4. On the **Storage Options** page, choose the appropriate storage type for your database files and then

click **Next**.

Note Ensure the storage is striped and mirrored. We recommend that you use Oracle ASM for managing database storage. Also, place datafiles, redo logs, and archived logs on separate disks to distribute I/O. For more information, refer to the I/O Configuration and Design section of the [Oracle Database Performance Tuning Guide](#).

5. On the **Recovery Configuration** page, choose the appropriate recovery options. For production environments, we recommend enabling flashback recovery and archive mode.

6. On the **Database Content** page, clear each of the additional components.

Note The Enterprise Manager Repository option can remain selected if you intend to use it.

7. On the **Initialization Parameters** page, on the **Memory** tab, use the following recommended settings as a starting point to configure the Perceptive Content Database memory.

Notes

- The settings should not exceed your server capacity, as paging may occur. If paging occurs, we recommend adding memory to the server or reducing Oracle memory settings.
- Depending on the configuration of the operating system, you may receive a performance boost by using huge pages for Oracle memory allocation. For more information on this, refer to the appropriate version-specific operating system and Oracle documentation.
- Use automatic shared memory management (ASMM), rather than automatic memory management (AMM).

Parameter	% of Recommended Memory	Example Value
SGA_MAX_SIZE	70% or less of total memory	20G
SGA_TARGET	*Same value as SGA_MAX_SIZE	20G
PGA_AGGREGATE_TARGET		4G

Note The values above are simply a starting point. After load testing with a sufficient representation of data and peak activity, review the AWR reports to determine the various advisory recommendations for each individual component of memory and adjust as necessary.

8. On the **Initialization Parameters** page, on the **Character Sets** tab, specify **WE8MSWIN1252** as the default character set for the INOW database.

Notes

- The National Character Set (AL16UTF16) can remain as the default value.
- Perceptive Content using an Oracle database does not currently support Unicode. If you require a Unicode implementation, we recommend that you consider using SQL Server or PostgreSQL.
- Since Oracle does not currently support Unicode, ensure that you define the **NLS_LANG** environment variable correctly to reflect the actual code page/character of the **client** before connecting to the database.

For example: NLS_LANG=AMERICAN_AMERICA.WE8MSWIN1252

The application server will automatically set the NLS_LANG environment variable on startup as shown in the example above. On Linux, the system sets the code page as follows within the connection string: IANAAppCodePage=2252

9. On the **Initialization Parameters** page, click **All Initialization Parameters**.
10. In the **All Initialization Parameters** dialog, click **Show Advanced Parameters**.

Note The value for each of the parameters in the dialog must match the values listed below.

Parameter	Value
QUERY_REWRITE_ENABLE	TRUE
QUERY_REWRITE_INTEGRITY	TRUSTED
CURSOR_SHARING	EXACT
NLS_DATE_FORMAT	RRRR-MM-DD
NLS_TIME_FORMAT	HH24:MI:SS.FF
NLS_TIMESTAMP_FORMAT	RRRR-MM-DD HH24:MI:SS.FF

11. Click **Close** and then click **Next**.
12. Click **Next** and continue to the **Creation Options** page.
13. Select the **Create Database** and the **Save as a Database Template** options.
14. Click **Finish**.

Result A summary of the database parameters appears.

15. Click **OK** to create the database. This may take some time, depending on your system.

Result When the database is complete, a summary appears with database log and connection information.

16. Click **OK**

Default Tablespaces

Once you create the Oracle database, you can then create the tablespaces and INUSER database user and objects. The downloaded Perceptive Content Server files include the **SunflowerORA.sql** file.

The INUSER schema utilizes two tablespaces by default and the database scripts will expect the following tablespaces to exist during object creation. You are free to use your preferred storage options (ASM, OMF, etc.) for the database files.

You can manually pre-create the tablespaces or you can update the SunflowerORA.sql file and define the datafile locations and the script will attempt to create the tablespaces for you. Be sure to configure datafiles to autoextend appropriately.

The **DATA** tablespace contains all the tables. If the script will be creating the tablespace then be sure to modify the path appropriately for the v_data_location variable.

The **INDX** tables contains all the indexes. If the script will be creating the tablespace then be sure to modify the path appropriately for the v_index_location variable.

About the INUSER Database Schema (User)

The INUSER database user is the user that owns all of the Perceptive Content database objects. This is the same user that Perceptive Content uses to connect to the INOW database.

The Perceptive Content database utilizes the NLS_COMP=LINGUISTIC and NLS_SORT=BINARY_CI parameters to implement case-insensitive searching, sorting and uniqueness within the database. This requires function-based indexes using the NLSSORT function to facilitate efficient index usage. The SunflowerORA.sql script will create all the necessary indexes.

The Perceptive Content Application Server will execute the following alter session commands for all database connections established by each of the service/agent connections pools.

```
ALTER SESSION SET NLS_TIMESTAMP_FORMAT = 'RRRR-MM-DD HH24:MI:SS.FF';
ALTER SESSION SET NLS_TIME_FORMAT = 'HH24:MI:SS.FF';
ALTER SESSION SET NLS_DATE_FORMAT = 'RRRR-MM-DD';
ALTER SESSION SET NLS_COMP=LINGUISTIC;
ALTER SESSION SET NLS_SORT=BINARY_CI;
ALTER SESSION SET OPTIMIZER_MODE = ALL_ROWS;
ALTER SESSION SET QUERY_REWRITE_INTEGRITY = TRUSTED;
ALTER SESSION SET QUERY_REWRITE_ENABLED = TRUE;
ALTER SESSION SET CURSOR_SHARING = EXACT;
ALTER SESSION SET OPTIMIZER_INDEX_COST_ADJ = 1;
```

Create the INUSER database user using SQL*Plus

The SunflowerORA.sql script will attempt to create the INUSER account if it does not already exist and if the default (DATA) and index (INDX) tablespaces exist. Alternatively, you can manually create the INUSER account.

You can create the INUSER database user using SQL*Plus or any other tool that allows you to execute SQL against the INOW database. Ensure that the INUSER account has the following minimum privileges. You can grant additional privileges to INUSER as needed. To create the INUSER database user using SQL*Plus, complete the following step.

Run the following commands:

```
create user INUSER identified by imagenow
default tablespace DATA
account unlock;

grant create session, resource, select any dictionary to inuser;

alter user inuser quota unlimited on DATA;
alter user inuser quota unlimited on INDX;
```

Create the database objects

1. Connect to the INOW database as SYS or another user that contains the Database Administrator role.

2. Execute the `SunflowerORA.sql` script.

This script will attempt to create the DATA and the INDX tablespaces as well as the INUSER and INEMUSER users if they do not already exist. It will also create the tables, indexes, and constraints.

Set up Oracle to use Secure Sockets Layer (SSL) security

SSL security uses cryptography and symmetric encryption to provide communication at the transport layer for data sent over a network for application-specific protocols. Mutual authentication, also known as two-way authentication, ensures that both servers trust each other. Please refer to the *SSL Security with Oracle Database* documentation for assistance with configuring two-way SSL authentication and communication between the Perceptive Content Database and Application Server.

Note You must have Oracle Enterprise Edition with the Advanced Security option installed on your database server. If you do not have the Oracle client installed on your application server, you can create the client server wallet and certificates on the database server and copy over the client server wallet

Create a PostgreSQL INOW database

The PostgreSQL database uses a Unicode (UTF8) encoding by default and is supported for Unicode implementations of Perceptive Content.

The INOW database uses the CITEXT extension in order to create several case-insensitive (CI) domains for character data types that are used across the inuser schema.

The **SunflowerPG.sql** file is included in the downloaded Perceptive Content Server files.

Locating the PostgreSQL configuration files

By default the `pg_hba.conf` and `postgresql.conf` files are located in the data directory of the database cluster (\$PGDATA).

To find the files, complete any of the following actions.

Command	Scope
<pre>show hba_file; show config_file; select name, setting from pg_settings where category = 'File Locations';</pre>	Connected to the database with the postgres role
<pre>\$ psql -U postgres -d postgres -c "show hba_file" \$ psql -U postgres -d postgres -c "show config_file"</pre>	Execute the command from the OS command line
Go to Tools then Server Configuration to view and change settings	pgAdmin III (Native GUI interface)

Configure PostgreSQL to allow remote connections

1. In the `pg_hba.conf` file, add a line for each server that will connect to the INOW database.

Note To restrict the establishment of connections, we recommend using a specific IP address. Although not recommended for security reasons, you have the ability to use “all” to allow connections from all servers.

The address specifies the set of hosts that matches the records. It can be a host name or a combination of an IP address and CIDR mask that contains an integer between 0 and 32, for IPv4; or between 0 and 128, for IPv6 that specifies the number of the significant bits in the mask.

2. After making updates, restart the PostgreSQL service “pg_ctl restart” or call “pg_ctl reload”.

This signals a SIGHUP to the postmaster and initializes the changes.

Note You must restart the PostgreSQL service; otherwise, the Perceptive Content Server will not be able to connect to the INOW database.

Configure connection settings

Update hba_file.conf

To find the file logged in as postgres, run the following command.

```
psql -U postgres -d postgres -c "show hba_file"
```

Add “all” to allow all, otherwise for access that is more restrictive specify a specific database, user and IP address for all connections to the database. For the address field append /32 for IPv4 addresses and /128 for IPv6 addresses.

# TYPE	DATABASE	USER	ADDRESS	METHOD
# IPv4 local connections:				
host	INOW	inuser	<<inserver ipaddress>>/32	md5
host	all	all	127.0.0.1/32	md5

Update postgresql.conf

To update postgresql.conf, modify the **listener_addresses** and **max_connections** parameters.

Note Changes to this file require a PostgreSQL Server restart

```
#-----
# CONNECTIONS AND AUTHENTICATION
#-----
# - Connection Settings -
listen_addresses = '*'          # what IP address(es) to listen on;
                                # comma-separated list of addresses;
                                # defaults to 'localhost'; use '*' for all
                                # (change requires restart)
port = 5432                    # (change requires restart)
max_connections = 300          # (change requires restart)
```

Default Tablespaces

The INOW database requires three tablespaces by default and the database scripts will expect the following tablespaces to exist during object creation.

The tablespaces can be manually pre-created or you can update the SunflowerPG.sql file and define the tablespace locations and the script will attempt to create the tablespaces for you.

If the script will be creating the tablespaces for you then be sure to modify the LOCATION appropriately within the script for each of the following three tablespaces.

Note The directories that you specify must already exist; They will not be created for you.

The **INOW_System** tablespace will contain all the postgres specific system objects.

The **INOW_Data** tablespace will contain all the inuser schema tables.

The **INOW_Indx** tablespace will contain all the inuser schema indexes.

Create the INOW database

To create the Perceptive Content INOW Database, complete the following steps.

1. Log in to the postgres database as the postgres user.
2. Execute the PostgreSQL version of the full DDL script: SunflowerPG.sql.

Note: You must execute the Perceptive Content DDL scripts from the command line using the PSQL utility. Specific psql commands within the script will not work with GUI tools, such as pgAdmin III.

```
EXAMPLE CONNECT STRING:
```

```
psql -h <IP_ADDRESS or localhost> -d postgres -U postgres -p 5432
```

```
EXAMPLE OF EXECUTING A SCRIPT WITHIN PSQL:
```

```
\i C:/DDL/7.2.0/ddl/SunflowerPG.sql  
\i 'C:/DDL/7.2.0/ddl/SunflowerPG.sql'
```

Database access through INEMUSER

In addition to the INUSER database user, you have the option to use the INEMUSER database user. The database setup script creates the INUSER and INEMUSER users. If your system includes applications that need to communicate with Perceptive Content Server, you can use the following database tables to process external messages. INEMUSER user can add, remove, and update records in the database tables without the risk of corrupting information stored in the Perceptive Content database. INEMUSER allows you to connect your external applications to the Perceptive Content database to access the tables. External applications, such as HL7 Interface, should only access the Perceptive Content database only through INEMUSER.

- IN_EXTERN_MSG
- IN_EXTERN_MSG_PROP
- IN_EXTERN_MSG_GROUP
- IN_EXTERN_MSG_SEQ

Order processing of external messages

Messages with the same group process in first-in-first-out (FIFO) order.

To group messages, complete the following steps.

1. In the IN_EXTERN_MSG_GROUP table, insert the EXTERN_MSG_ID and an EXTERN_MSG_GROUP_ID.
2. In the IN_EXTERN_MSG_SEQ table, insert the EXTERN_MSG_ID.

Notes

- The EXTERN_MSG_ID is the same ID you used in the IN_EXTERN_MSG table.
- You can add more than one group to a message by inserting additional records into the IN_EXTERN_MSG_GROUP table.
- The EXTERN_MSG_ID and EXTERN_MSG_GROUP_ID columns have a limit of 64 characters.
- To create a new message, insert a record into the IN_EXTERN_MSG table with a unique value for the EXTERN_MSG_ID column.

Unlock the INEMUSER account

The INEMUSER account is locked by default. You must unlock the account if you intend to utilize it.

To unlock the INEMUSER account, perform one of the following actions according to the RDBMS you are using.

- For SQL Server, execute the following command.

```
alter login inemuser enable;
```

- For Oracle, execute the following command.

```
alter user inemuser account unlock;
```

- For PostgreSQL, execute either of the following commands based on the preferred active date range.

```
alter role inemuser valid until 'YYYY-MM-DD 00:00:00';
```

```
alter role inemuser valid until 'infinity';
```