

openCRX Installation Guide for PostgreSQL 8

Version 2.4



www.opencrx.org

License

The contents of this file are subject to a BSD license (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at <http://www.opencrx.org/license.htm>

Copyright 2009 © CRIXP Corp. All rights reserved.



Table of Contents

1	About this Book.....	4
1.1	Who this book is for.....	4
1.2	What do you need to understand this book.....	4
1.3	Tips, Warnings, etc.....	4
2	Prerequisites.....	5
2.1	Issue(s) specific to PostgreSQL.....	6
2.1.1	Puzzling behavior of PostgreSQL.....	6
2.1.2	Implications for openCRX.....	6
2.1.3	Enhancing performance.....	7
2.1.4	PostgreSQL Installation Hints / Locale Support.....	8
3	Upgrading from previous versions.....	9
3.1	The SQL Script upgrade-from-.....	9
3.2	The SQL Script migrate-from-.....	9
3.3	The SQL Script drop-from-.....	9
3.4	The SQL Script dbcreate-views.sql.....	10
3.5	The SQL Script dbcreate-indexes.sql.....	10
3.6	Populate Preferences.....	10
4	Create the database.....	11
4.1	Create the database with psql terminal.....	11
4.2	Create the database with pgAdmin III.....	13
5	Install the openCRX Database Schema Objects.....	16
5.1	Install database schema objects with psql terminal.....	16
5.2	Install database schema objects with pgAdmin III.....	17
6	Next Steps.....	19

List of Figures

Figure 1: Initdb parameters of pgInstaller on Windows.....	8
Figure 2: Start pgAdmin III and connect to the database.....	13
Figure 3: Create a new Login Role system – Step 1.....	13
Figure 4: Create a new Login Role system – Step 2.....	14
Figure 5: Create a new Database – Step 1.....	14
Figure 6: Create a new Database – Step 2.....	15
Figure 7: Create a new Database – Step 3.....	15
Figure 8: Start pgAdmin III Query.....	17
Figure 9: Load and execute dbcreate-tables.sql in pgAdmin III Query.....	17

List of Listings

Listing 1: Configure access to DB.....	11
Listing 2: Remove ident sameuser.....	11
Listing 3: Stop PostgreSQL server.....	11
Listing 4: Start PostgreSQL server.....	11
Listing 5: Create user system with psql terminal.....	12
Listing 6: Create the database CRX_CRX with psql terminal.....	12
Listing 7: Install database schema objects with psql terminal.....	16

1 About this Book

This book describes how to setup an openCRX database instance for PostgreSQL.

1.1 Who this book is for

The intended audience are openCRX database administrators.

1.2 What do you need to understand this book

This book describes the installation of openCRX for PostgreSQL. The book assumes that you are familiar with PostgreSQL installation and configuration.

1.3 Tips, Warnings, etc.

We make use the following pictograms:



Information provided as a “Tip” might be helpful for various reasons: time savings, risk reduction, etc.



You should carefully read information marked with “Important”. Ignoring such information is typically not a good idea.



Warnings should not be ignored (risk of data loss, etc.)

2 Prerequisites

As a first step you must download the following software packages:

- Download and install **openCRX SDK Installer**. It is available from <http://www.opencrx.org/sdk.htm>
The SDK contains the DB scripts required to install an *openCRX* database in the directory `<SDK_Install_Dir>\opencrx-x.x.x\core\src\sql`
- Download **PostgreSQL Database Server** from <http://www.postgresql.org/download/>



We recommend PostgreSQL v8.3 as the performance is much better than with previous versions.

For Windows, get the **pgInstaller** (and not the “one click installer”) – the reason is explained in chapter 2.1.4 PostgreSQL Installation Hints / Locale Support.

- Download **pgAdmin III** from <http://www.postgresql.org/download/>
- Download the current **PostgreSQL JDBC 3 driver** from <http://jdbc.postgresql.org/download.html>
The JDBC driver is required for the application server installation.



Please ensure that you install the **correct JDBC driver** (i.e. correct in terms of type, version, etc.) and **one JDBC driver** only! Ignoring this wisdom leads to problems as the connection to the database will fail.

Read the information provided at <http://jdbc.postgresql.org/> carefully as picking the wrong driver will cause you misery. For example, with PostgreSQL v8.3 and openCRX v2.2 (and hence JDK 1.5) you should choose the current JDBC 3 driver “8.3-604 JDBC 3” (as of 7-MAY-2009).



Read the next chapter “Issue(s) specific to PostgreSQL” **before** you install PostgreSQL – at a minimum you should read chapter 2.1.4 PostgreSQL Installation Hints / Locale Support to understand the importance of a correctly configured database cluster.

As a next step you must install **PostgreSQL** and **pgAdmin III** (please refer to the PostgreSQL documentation for installation details).

2.1 Issue(s) specific to PostgreSQL

Like any other DBMS, PostgreSQL has some issues and we want you to be aware of them.

2.1.1 Puzzling behavior of PostgreSQL

Based on our analysis, it seems PostgreSQL is not behaving consistently across platforms. The replies to the following (trivial) select statements are absolutely irritating:

Select Statement	expected reply	pg 8.3 on Windows	pg 8.3 on Kubuntu 7.0.4	pg 8.2 on Kubuntu 7.0.4
<code>select '0' > '/'</code>	true	true	true	true
<code>select '0' > '/a'</code>	true	true	false	false
<code>select '0' > '/aa'</code>	true	true	false	false

The replies of pg 8.3 on Windows are correct, **some of the replies of pg on Linux are (in our opinion) not correct.** This puzzling behavior is the reason why you might need some **special settings to improve performance of openCRX with PostgreSQL.**

A note to the PostgreSQL community:

We are aware of locale-specific sorting, but the results of the above 3 select statements **should be true for any locale** as

'0' > '/' implies '0' > '/(.)+'

(unless '/' is treated as some kind of special escape character in a particular locale so that '/a < '0' is true for such a locale; to our best knowledge, such a locale does not exist).



Details of PostgreSQL locale support are available from <http://www.postgresql.org/docs/8.3/interactive/locale.html>

2.1.2 Implications for openCRX

As "object ID matching" (OID matching) is a frequent operation it is absolutely crucial that it can be done in a very efficient way, otherwise openCRX will suffer from a heavy performance hit. The openCRX database plugin does OID matching with SQL statements containing comparisons like

(object_id > id_pattern_0) and (object_id < id_pattern_1)

Given the issues that exist with PostgreSQL the default configuration of the openCRX database plugin resorts to a comparison based on LIKE. We are aware of the implications – a severe performance hit – as prepared statements with LIKE comparisons typically don't use indices.

2.1.3 Enhancing performance

If any of the following conditions is true you can override the default setting of the openCRX database plugin:

- the locale of the openCRX database is equal to C
- the locale of the openCRX database is equal to POSIX
- **lc_collate** and **lc_ctype** of your pg cluster are equal to C; open a query window in pgAdmin and execute "show all" to see the settings of your database cluster:

	name text	setting text	description text
89	lc_collate	C	Shows the collation order locale.
90	lc_ctype	C	Shows the character classification and case conversion locale.

With PostgreSQL, the system property

`org.openmdx.persistence.jdbc.useLikeForOidMatching`

is by default set to true; this ensures that OID matching works as expected, but the price is a severe performance hit. If any of the above conditions is satisfied you can safely set this system property to false resulting in much improved performance. You can override the default setting by providing the following startup option to openCRX:

`-Dorg.openmdx.persistence.jdbc.useLikeForOidMatching=false`



Setting `useLikeForOidMatching` to false without ensuring that

- either **the database locale is equal to C or POSIX**
- or **lc_collate and lc_ctype are equal to C**

is asking for severe trouble as OID matching will not work as intended. Sooner or later you will run into data consistency issues.

Hence, initialize your database cluster as described below or leave `useLikeForOidMatching` at its default value (true).

2.1.4 PostgreSQL Installation Hints / Locale Support

Assuming you want full UTF-8 support for the openCRX database, we recommend you set the locale behavior of the server to C as follows.

On **Linux**,

- install the latest version of PostgreSQL (and the required libraries)
- install pgadmin
- please note that the package installer runs `initdb` automatically, but probably not with the desired locale settings (by default, `initdb` is run with the same locale as your OS default locale, i.e. typically UTF-8; this leads to problems as indicated in chapter 2.1.1 Puzzling behavior of PostgreSQL), hence we need to redo the initialization:
 - `stop pg daemon`
 - `rename /var/lib/pgsql/data to /var/lib/pgsql/data.ori`
 - `su - postgres`
 - `initdb --pgdata=data --encoding=UTF8 --locale=C --username=postgres --pwprompt`

Alternatively, you can set "LC_ALL=C" in the environment of the relevant pg processes; if LC_ALL is set to "C", the PostgreSQL sort order is also correct.

On **Windows**, use the **pgInstaller** which allows you to set `initdb` parameters as follows (more information at <http://pginstaller.projects.postgresql.org/>):

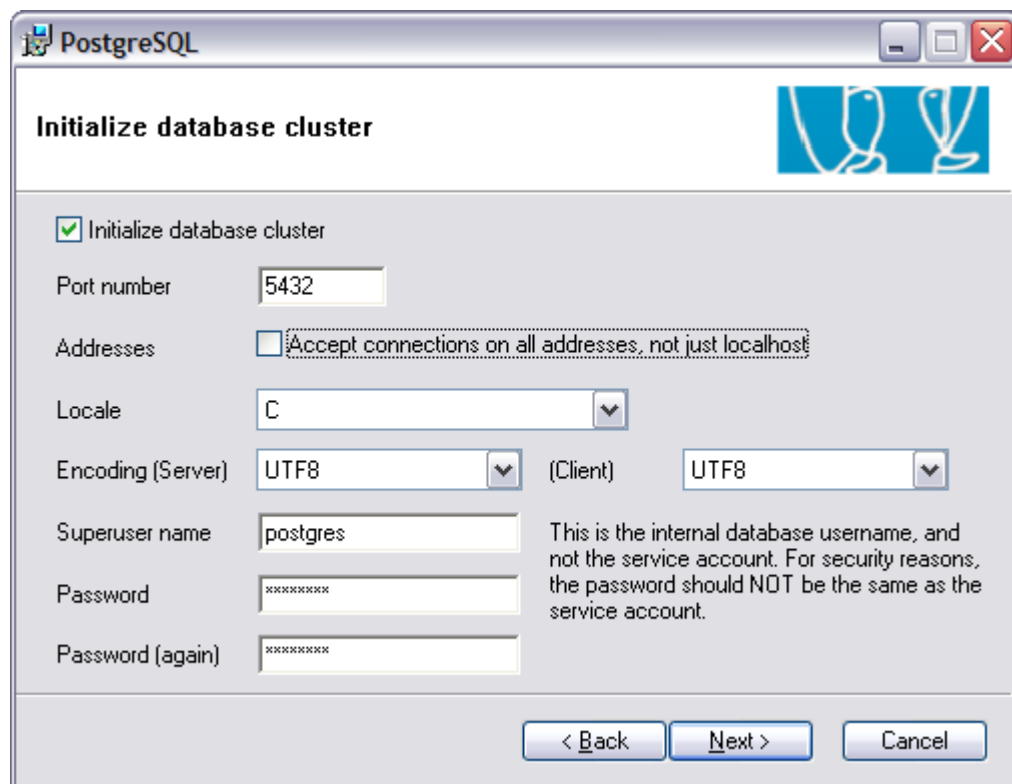


Figure 1: *Initdb parameters of pgInstaller on Windows*

Warning: Note that the **one click installer does not set the locale to C!**

3 Upgrading from previous versions

If you already have PostgreSQL for openCRX installed, upgrade the database as explained in this chapter. You can then skip the rest of the document.



Warning

Backup your database **before** you run any DB scripts!



Warning

Please consult <http://www.opencrx.org/faq.htm#upgrade> and find out whether there exist specific instructions for your openCRX version. Instructions below are generic and might not cover all steps required to successfully upgrade your openCRX version.



Warning

Please note that you cannot skip openCRX versions when upgrading your database.

3.1 The SQL Script upgrade-from-...

In a first step you must upgrade your database. openCRX distributions provide an SQL script of the form

upgrade-from-<version from>-to-<version to>.sql

If you have installed openCRX 2.3.0, for example, and you want to upgrade to version 2.4.0 you have to run the script `upgrade-from-2.3.0-to-2.4.0.sql` on your database instance.

3.2 The SQL Script migrate-from-...

In a second step you must migrate your database. openCRX distributions often times provide an SQL script of the form

migrate-from-<version from>-to-<version to>.sql

If you have installed openCRX 2.3.0, for example, and you want to upgrade to version 2.4.0 you have to run the script `upgrade-from-2.3.0-to-2.4.0.sql` on your database instance.

3.3 The SQL Script drop-from-...

Next you can drop unused tables from your database. openCRX distributions often times provide an SQL script of the form

drop-from-<version from>-to-<version to>.sql

If you have installed openCRX 2.3.0, for example, and you want to drop tables not used by openCRX 2.3.0 you can run the script `drop-from-2.3.0-to-2.4.0.sql`

on your database instance. Alternatively, you can also rename such tables, e.g. from `transition_type` to `unused_transition_type`. Also, it goes without saying that you should never drop a table before you made a backup!

3.4 The SQL Script `dbcreate-views.sql`

Most new openCRX versions make use of new/changed views, i.e. if an openCRX distribution includes an SQL script of the form

`dbcreate-views.sql`

then you should run that script. If you have installed openCRX 2.3.0, for example, and you want to upgrade to openCRX 2.4.0 you should run the script `dbcreate-views.sql` on your database instance. Make sure that old views are indeed dropped and new views properly created.



If the `dbcreate-views.sql` script does not run without error you can proceed as follows:

- drop all the views manually
- remove all the **drop view** statements (located at the beginning) from the script `dbcreate-views.sql`
- execute the remainder of the script `dbcreate-views.sql` that contains **create view** statements only; these statements must run without error

3.5 The SQL Script `dbcreate-indexes.sql`

Most new openCRX versions make use of new/changed indexes, i.e. if an openCRX distribution includes an SQL script of the form

`dbcreate-indexes.sql`

then you should run that script. If you have installed openCRX 2.3.0, for example, and you want to upgrade to openCRX 2.4.0 you should run the script `dbcreate-indexes.sql` on your database instance.

3.6 Populate Preferences

The last step involves deleting old preferences and populating the table with new ones. Run the SQL script `populate-preferences.sql` to do this.



Make sure that old preferences are indeed removed and new ones loaded. This table contains the configuration of the openMDX database plugin, i.e. openCRX persistency will not work properly if the loaded preferences do not match the version of openCRX.

4 Create the database

You can either create the database with the psql terminal or with pgAdmin III.

4.1 Create the database with psql terminal

In order to connect to a PostgreSQL server from a remote pgAdmin III instance you have to properly configure the PostgreSQL server.

Let the PostgreSQL server accept connections from inside your network (suppose 192.168.1.0/24) by postgres and system users. We suppose the connections will be clear (not ssl) and protected by a password authentication mechanism. To do this, add the following lines at the end of the file `<postgres home dir>/data/pg_hba.conf`:

Listing 1: Configure access to DB

#TYPE	DATABASE	USER	IP-ADDRESS	IP-MASK	METHOD
local	all	all			password
host	CRX_CRX	system	192.168.1.0	255.255.255.0	password
host	all	postgres	192.168.1.0	255.255.255.0	password



Remove the line with ident sameuser if it exists in `pg_hba.conf`:

Listing 2: Remove ident sameuser

```
host all all 127.0.0.1 255.255.255.255 ident sameuser
```

If the PostgreSQL server is running, stop it before you continue:

Listing 3: Stop PostgreSQL server

```
postgres$pg_ctl stop
```

Next, the PostgreSQL server has to be configured to accept more than one connection through a socket. To do this, log-in to the system with the PostgreSQL server user account (usually `postgres`) and start the PostgreSQL server with the following options:

Listing 4: Start PostgreSQL server

```
postgres$pg_ctl -o -i -l postmaster.log start
```

Now you can connect to the default PostgreSQL server database (template1) with the `postgres` user (DB Administrator).

Create a PostgreSQL *user* named `system`. Set the password of the user to `manager` (for the purpose of this guide). Allow the user to create databases.

Listing 5: Create user system with psql terminal

```
su - postgres          # change to your postgresql Account.
createuser -P system   # Create the user system and ask for a password
```

Create the database `CRX_CRX` with the following commands:

Listing 6: Create the database CRX_CRX with psql terminal

```
su - postgres
createdb -h localhost -E utf8 -U system CRX_CRX
```



Set the Encoding to UTF-8 if you intend to make use of the openCRX UTF-8 support (but make sure you fully understand the information provided in chapters 2.1.1 Puzzling behavior of PostgreSQL and 2.1.4 PostgreSQL Installation Hints / Locale Support.



Please note that you are free to choose any database name you like, i.e. there is no requirement to name the database `CRX_CRX`.

However, you must ensure that the database name you choose is also used in the respective **database definition file** (e.g. `openejb.xml` for Tomcat) and any other configuration that contains the database name.

For example, don't forget to adapt the following Java option `-Dopencrx.CRX.jdbc.url="jdbc:postgresql://localhost:5432/CRX_CRX"` accordingly.

You have completed creating the openCRX database.

4.2 Create the database with pgAdmin III

Start pgAdmin III and connect to the appropriate database server:

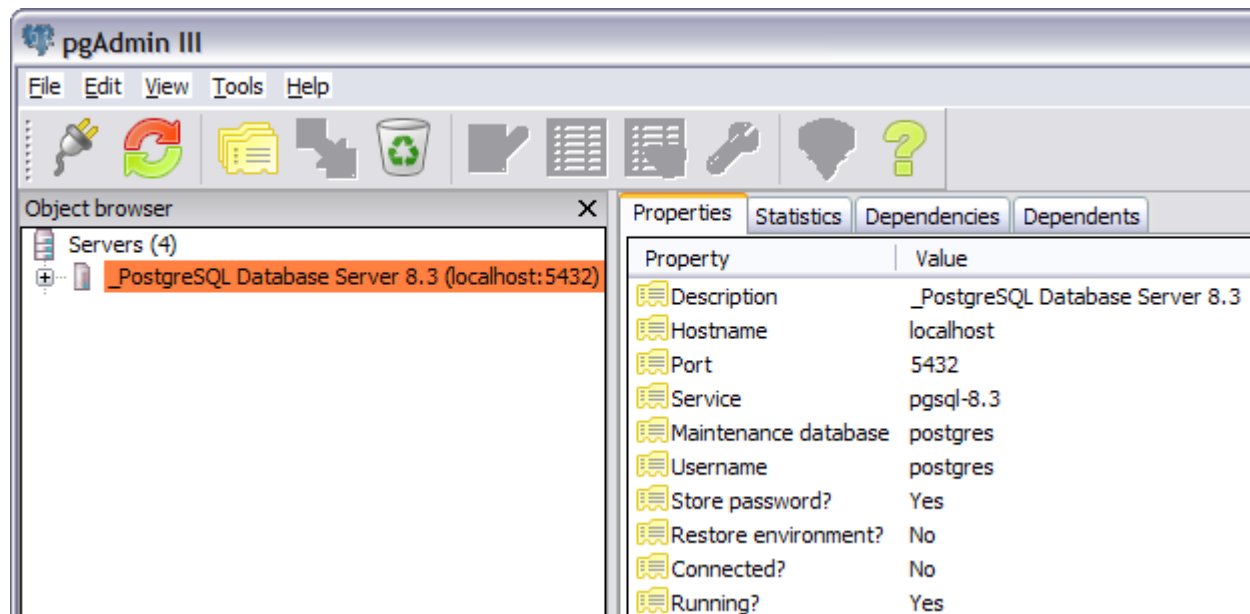


Figure 2: Start pgAdmin III and connect to the database

Next you create a new *login role* **system** as follows:

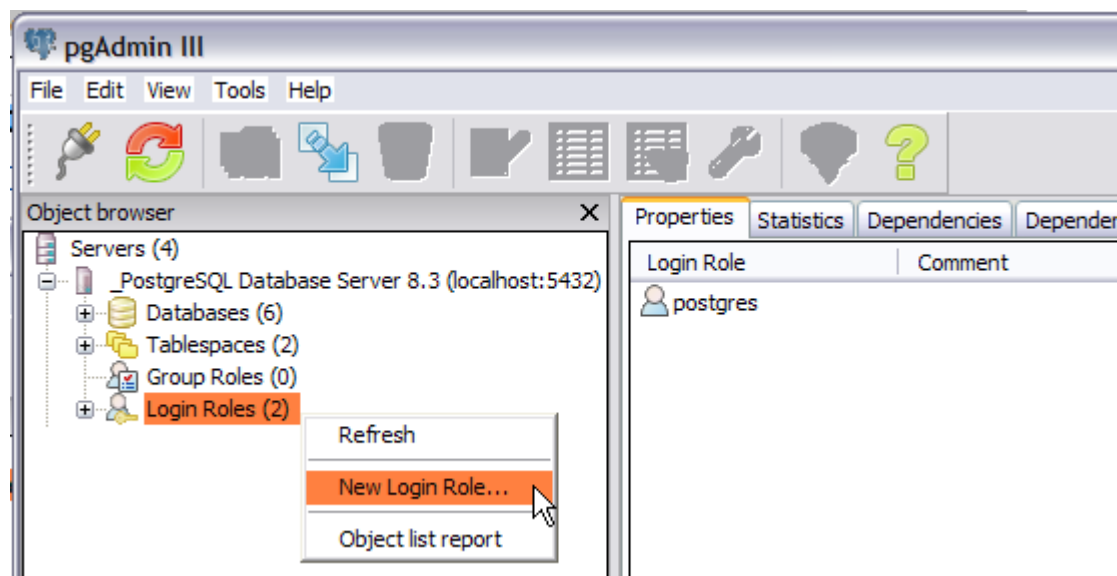


Figure 3: Create a new Login Role system – Step 1

You will get a new window **New Login Role** where you can enter the *Role name* **system** and a *password* (we use *password manager* for the purpose of this guide) – verify that you check options as shown below:

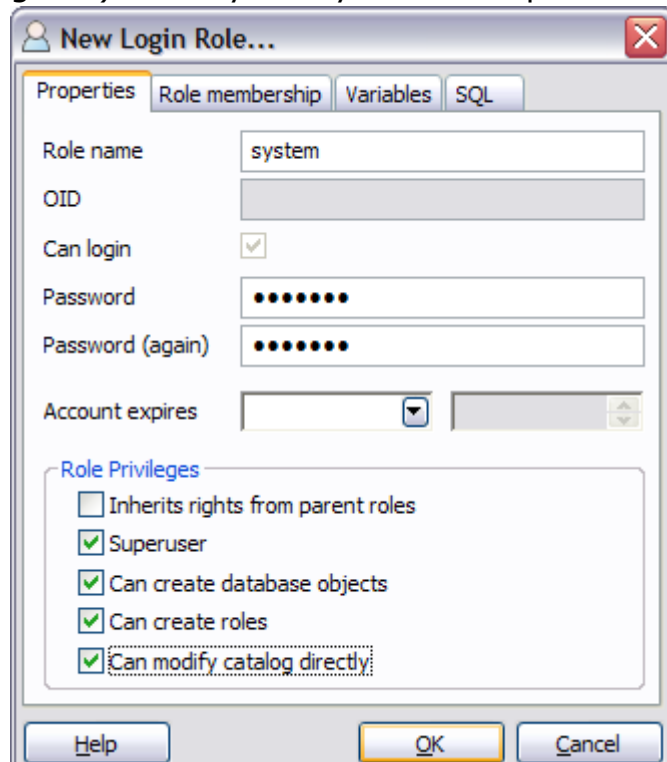


Figure 4: Create a new Login Role **system** – Step 2



Please note that the login role **system** must be a **Superuser**

pgAdmin III allows you to create and manage databases. Before you can install the openCRX database schema objects you must create a new database for openCRX. Right-click on the tree item **Databases** and then select the pop-up menu entry **New Database** as shown below:

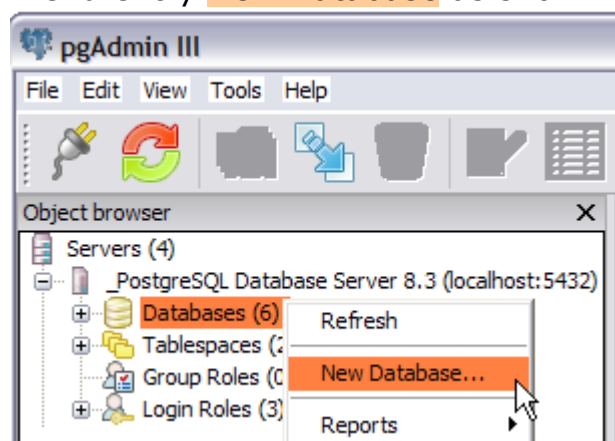
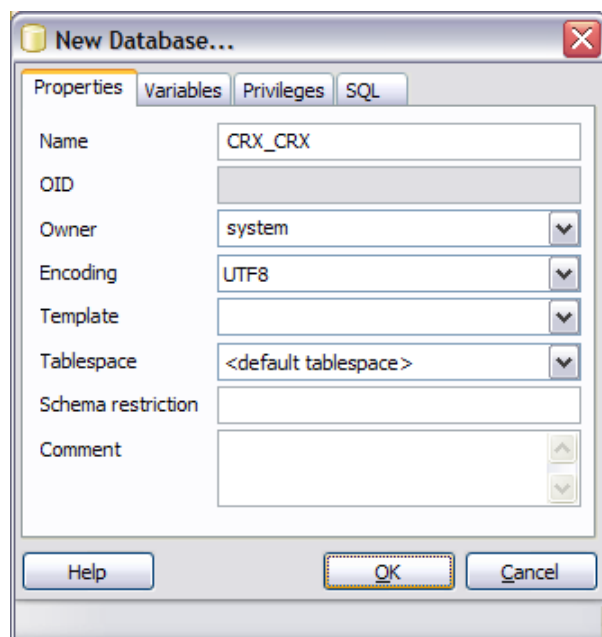


Figure 5: Create a new Database – Step 1

You will get a new window New Database – populate it as shown below to create the database **CRX_CRX** owned by **system**:



It is highly recommended to set the encoding to **UTF-8** so that you can make use of the openCRX's multi-language support that requires a UTF-8 encoded database.

Figure 6: Create a new Database – Step 2

After creating the new Database **CRX_CRX** your pgAdmin III window should look similar to the following figure:




Figure 7: Create a new Database – Step 3

You have completed creating the database **CRX_CRX**.

5 Install the openCRX Database Schema Objects

After creating the schema you are now ready to install the openCRX database schema objects. The following scripts must be executed in the order given below:

- dbcreate-tables.sql
- dbcreate-views.sql
- dbcreate-indexes.sql
- populate-preferences.sql



Do not execute any other scripts included in the distribution.

Again, depending on your preferences you can either use the psql terminal or pgAdmin III to install the openCRX database schema objects.

5.1 Install database schema objects with psql terminal

Execute the scripts in the given order.



If any of the scripts does not run without errors, correct the errors **before** you continue with the next script.

Listing 7: Install database schema objects with psql terminal

```
psql -U system CRX_CRX < dbcreate-tables.sql
psql -U system CRX_CRX < dbcreate-views.sql
psql -U system CRX_CRX < dbcreate-indexes.sql
psql -U system CRX_CRX < populate-preferences.sql
```



The script **dbcreate-views.sql** tries to drop old views before it creates the new ones. In case there are no existing views you need to comment out the drop statement to successfully create new views.

This completes the installation of the openCRX database schema objects.

5.2 Install database schema objects with pgAdmin III

After creating the database instance and the user `system` you are ready to install the openCRX database schema objects. Connect to the database `CRX_CRX` as user `system` and then start **pgAdmin III Query** as shown below:





Figure 8: Start pgAdmin III Query

Next you load the script `dbcreate-tables.sql` into the query window and execute it by clicking on the play button :

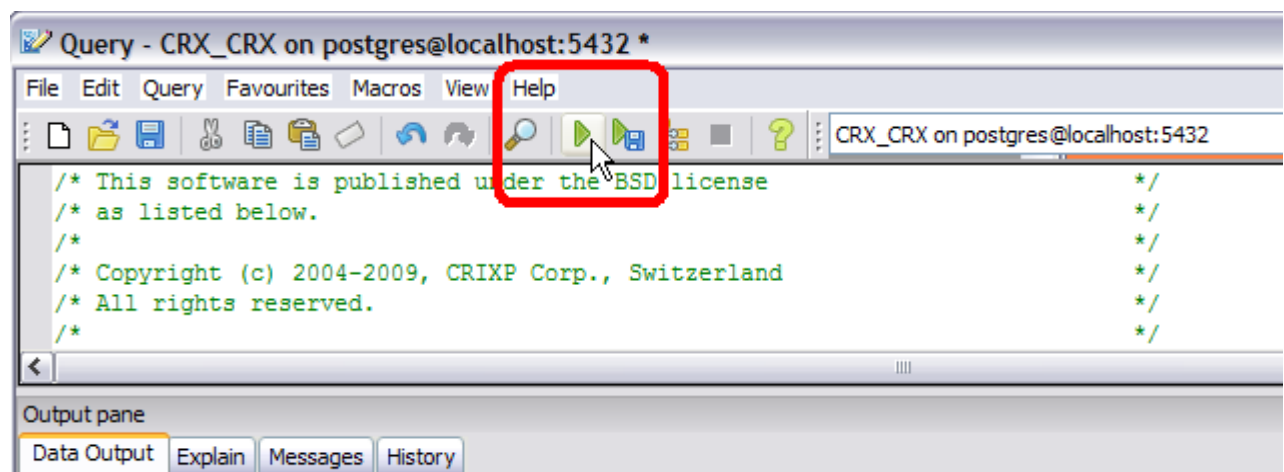


Figure 9: Load and execute `dbcreate-tables.sql` in pgAdmin III Query



If any of the scripts does not run without errors, correct the errors **before** you continue with the next script.

Similarly, execute the remaining scripts in the following order:

- **dbcreate-views.sql**
- **dbcreate-indexes.sql**
- **populate-preferences.sql**



The script **dbcreate-views.sql** tries to drop old views before it creates the new ones. In case there are no existing views you need to comment out the drop statement to successfully create new views.

The scripts should run without errors and after execution you might want to verify with pgAdmin III that all the objects (i.e. tables, views, indexes, and a populated table `prefs_preference`) were properly created.

This completes the installation of the openCRX database schema objects.

6 Next Steps

If you have completed successfully the database installation you are ready to use the openCRX database **CRX_CRX**. The application server installation guides explain how to connect the application server to the openCRX database instance.