

# « E-Maj » PostgreSQL extension

-

## User's Guide

Version: 0.10.1

# Summary

|  |           |
|--|-----------|
| <b>1 - <a href="#">Introduction</a></b> .....                                    | <b>5</b>  |
| 1.1 - <a href="#">Document content</a> .....                                     | 5         |
| 1.2 - <a href="#">License</a> .....  | 5         |
| 1.3 - <a href="#">E-Maj's objectives</a> .....                                   | 5         |
| <b>2 - <a href="#">How E-Maj works</a></b> .....                                 | <b>7</b>  |
| 2.1 - <a href="#">Concepts</a> .....   | 7         |
| 2.1.1 - <a href="#">Tables Group</a> .....                                       | 7         |
| 2.1.2 - <a href="#">Mark</a> .....   | 7         |
| 2.1.3 - <a href="#">Rollback</a> .....   | 7         |
| 2.2 - <a href="#">Architecture</a> .....   | 8         |
| 2.2.1 - <a href="#">Created objects</a> .....                                    | 8         |
| 2.2.2 - <a href="#">Norm for E-Maj objects naming</a> .....                      | 9         |
| <b>3 - <a href="#">How to install E-Maj</a></b> .....                            | <b>10</b> |
| 3.1 - <a href="#">extension download and decompression</a> .....                 | 10        |
| 3.1.1 - <a href="#">Download</a> .....   | 10        |
| 3.1.2 - <a href="#">Decompression</a> .....                                      | 10        |
| 3.2 - <a href="#">E-Maj extension setup</a> .....                                | 11        |
| 3.2.1 - <a href="#">Full installation</a> .....                                  | 11        |
| 3.2.2 - <a href="#">Update an existing E-Maj version prior 0.10.0</a> .....      | 13        |
| 3.2.3 - <a href="#">Migration from an already installed 0.10.0 version</a> ..... | 14        |
| 3.2.4 - <a href="#">Changes in postgresql.conf configuration file</a> .....      | 15        |
| 3.2.5 - <a href="#">Test and demonstration</a> .....                             | 15        |
| 3.3 - <a href="#">E-Maj uninstall</a> .....                                      | 16        |
| <b>4 - <a href="#">How to use E-Maj</a></b> .....                                | <b>18</b> |
| 4.1 - <a href="#">Set-up the E-Maj access policy</a> .....                       | 18        |
| 4.1.1 - <a href="#">E-Maj roles</a> .....  | 18        |
| 4.1.2 - <a href="#">Giving E-Maj rights</a> .....                                | 18        |
| 4.1.3 - <a href="#">Giving rights on application tables and objects</a> .....    | 19        |
| 4.1.4 - <a href="#">Synthesis</a> .....  | 19        |
| 4.2 - <a href="#">Main functions</a> .....                                       | 20        |
| 4.2.1 - <a href="#">Operations chain</a> .....                                   | 20        |
| 4.2.2 - <a href="#">Define tables groups</a> .....                               | 21        |

|  |    |
|--|----|
| 4.2.3 - <a href="#">Create a tables group</a> .....                            | 22 |
| 4.2.4 - <a href="#">Start a tables group</a> .....                             | 22 |
| 4.2.5 - <a href="#">Set an intermediate mark</a> .....                         | 23 |
| 4.2.6 - <a href="#">Rollback a tables group</a> .....                          | 24 |
| 4.2.7 - <a href="#">Perform a logged rollback of a tables group</a> .....      | 25 |
| 4.2.8 - <a href="#">Stop a tables group</a> .....                              | 26 |
| 4.2.9 - <a href="#">Drop a tables group</a> .....                              | 27 |
| 4.2.10 - <a href="#">Change a tables group content</a> .....                   | 27 |
| 4.2.11 - <a href="#">Change the structure of an application table</a> .....    | 28 |
| 4.3 - <a href="#">Multi-groups functions</a> .....                             | 29 |
| 4.3.1 - <a href="#">General information</a> .....                              | 29 |
| 4.3.2 - <a href="#">Functions list</a> .....                                   | 29 |
| 4.3.3 - <a href="#">Syntax for groups array</a> .....                          | 29 |
| 4.3.4 - <a href="#">Other considerations</a> .....                             | 30 |
| 4.4 - <a href="#">Marks management functions</a> .....                         | 31 |
| 4.4.1 - <a href="#">Comments on marks</a> .....                                | 31 |
| 4.4.2 - <a href="#">Rename a mark</a> .....                                    | 31 |
| 4.4.3 - <a href="#">Delete a mark</a> .....                                    | 32 |
| 4.4.4 - <a href="#">Delete oldest marks</a> .....                              | 32 |
| 4.5 - <a href="#">Statistics functions</a> .....                               | 33 |
| 4.5.1 - <a href="#">Global statistics about logs</a> .....                     | 33 |
| 4.5.2 - <a href="#">Detailed statistics about logs</a> .....                   | 34 |
| 4.5.3 - <a href="#">Estimate the rollback duration</a> .....                   | 35 |
| 4.5.4 - <a href="#">Get a mark name</a> .....                                  | 36 |
| 4.6 - <a href="#">Other functions</a> .....                                    | 37 |
| 4.6.1 - <a href="#">Simultaneous rollback and stop of a tables group</a> ..... | 37 |
| 4.6.2 - <a href="#">Reset log tables of a group</a> .....                      | 37 |
| 4.6.3 - <a href="#">Comments on groups</a> .....                               | 37 |
| 4.6.4 - <a href="#">Check the consistency of the E-Maj environment</a> .....   | 38 |
| 4.6.5 - <a href="#">Forced suppression of a tables group</a> .....             | 39 |
| 4.6.6 - <a href="#">Snap tables of a group</a> .....                           | 39 |
| 4.6.7 - <a href="#">Snap log tables of a group</a> .....                       | 40 |
| 4.7 - <a href="#">Parallel Rollback</a> .....                                  | 42 |
| 4.7.1 - <a href="#">Sessions</a> .....   | 42 |

|  |           |
|--|-----------|
| 4.7.2 - <a href="#">Prerequisites</a> .....                                | 42        |
| 4.7.3 - <a href="#">Syntax</a> .....                                       | 43        |
| 4.7.4 - <a href="#">Examples</a> .....                                     | 44        |
| <b>5 - <a href="#">Miscellaneous</a></b> .....                             | <b>45</b> |
| 5.1 - <a href="#">Parameters</a> .....                                     | 45        |
| 5.2 - <a href="#">Checks</a> .....   | 46        |
| 5.3 - <a href="#">Traces of operations</a> .....                           | 47        |
| 5.4 - <a href="#">Impacts on cluster and database administration</a> ..... | 48        |
| 5.4.1 - <a href="#">Stopping and restarting the cluster</a> .....          | 48        |
| 5.4.2 - <a href="#">Saving and restoring the database</a> .....            | 48        |
| 5.4.3 - <a href="#">Tables reorganisation</a> .....                        | 50        |
| 5.4.4 - <a href="#">Using E-Maj with replication</a> .....                 | 50        |
| 5.4.5 - <a href="#">PostgreSQL version upgrade</a> .....                   | 50        |
| 5.5 - <a href="#">Sensitivity to system time change</a> .....              | 51        |
| 5.6 - <a href="#">Usage limits</a> .....                                   | 51        |
| 5.7 - <a href="#">User's responsibility</a> .....                          | 53        |
| 5.7.1 - <a href="#">Defining tables groups content</a> .....               | 53        |
| 5.7.2 - <a href="#">Appropriate call of main functions</a> .....           | 53        |
| 5.7.3 - <a href="#">Management of application triggers</a> .....           | 53        |
| 5.7.4 - <a href="#">Internal E-Maj table or sequence change</a> .....      | 54        |
| <b>6 - <a href="#">phpPgAdmin plugin</a></b> .....                         | <b>55</b> |
| 6.1 - <a href="#">General presentation</a> .....                           | 55        |
| 6.2 - <a href="#">Using phpPgAdmin plugin</a> .....                        | 55        |
| 6.2.1 - <a href="#">How to reach E-Maj from phpPgAdmin interface</a> ..... | 55        |
| 6.2.2 - <a href="#">Tables groups list</a> .....                           | 56        |
| 6.2.3 - <a href="#">Tables group composition</a> .....                     | 58        |
| 6.2.4 - <a href="#">Tables group details</a> .....                         | 59        |
| 6.2.5 - <a href="#">Rollback statistics</a> .....                          | 61        |
| 6.2.6 - <a href="#">Statistics about log tables content</a> .....          | 61        |

# 1 INTRODUCTION

---

## 1.1 DOCUMENT CONTENT

---

This document is a user's guide for the E-Maj PostgreSQL extension.

Chapter 2 presents the concepts used by E-Maj and the general architecture of the extension.

Chapter 3 describes in detail how to setup and use E-Maj.

Chapter 4 gives some additional information needed for a good understanding of how the extension works.

Then, chapter 5 presents the E-Maj extension of the administration tool phpPgAdmin.

## 1.2 LICENSE

---

This extension and its documentation are distributed under GPL license (GNU - General Public License).

## 1.3 E-MAJ'S OBJECTIVES

---

E-Maj is the French acronym for « Enregistrement des Mises A Jour », which means « updates recording ».

The main goal of E-Maj is to supply a way to logically restore sets of tables into predefined states, without being obliged to either restore all files of the PostgreSQL instance (cluster) or reload the entire content of the concerned tables.

But E-Maj may also be used to trace updates performed by application programs on the table's content .

It provides a good solution to :

- define save points at precise time on a tables group,
- restore if needed this tables group into a stable state, without stopping the cluster,
- manage several save points, each of them being usable at any time as restore point.

So, in a production environment, E-Maj may simplify the technical architecture, by offering a smooth and efficient alternative to time and/or disk consuming intermediate saves

(*pg\_dump*, *mirror disks*,...). E-Maj may also bring a help to the debugging by giving a way to precisely analyse how suspicious programs update application tables.

In a test environment, E-Maj also brings smoothness into operations. It is possible to very easily restore databases into predefined stable states, so that tests can be replayed as many times as needed.

## 2 HOW E-MAJ WORKS

---

### 2.1 CONCEPTS

---

E-Maj is built on three main concepts.

#### 2.1.1 Tables Group

The « *tables group* » represents a set of application tables that live at the same rhythm, meaning that their content can be restored as a whole if needed. Typically, it deals with all tables that are updated by one or several processings. Each tables group is defined by a name which must be unique inside its database. By extent, a tables group can also contain application sequences (in the RDBMS sense). Tables and sequences that constitute a tables group can belong to different schemas of the database.

At a given time, a tables group is either in a « logging » state or in a « *idle* » state. The logging state means that all updates applied on the tables of the group are recorded.

A tables group can be either “*rollback-able*”, which is the standard case, or “*audit\_only*”. In this later case, it is not possible to rollback the group. But with this type of group, it is possible to record tables updates for auditing purpose, even with tables that do not have primary key known in PostgreSQL catalogue.

#### 2.1.2 Mark

A « *mark* » is a particular point in the life of a tables group, corresponding to a stable point for all tables and sequences of the group. A mark is explicitly set by a user operation. It is defined by a name that must be unique for the tables group.

#### 2.1.3 Rollback

The « *rollback* » operation consists in resetting all tables and sequences of a group in the state they had when a mark has been set.

Actually, there are two rollback types:

- with a « *unlogged rollback* », no trace of updates that are cancelled by the rollback operation are kept,
- with « *logged rollback* », updates cancellations are recorded in log tables, so that they can be later cancelled: the rollback operation can be ... rolledback.

## 2.2 ARCHITECTURE

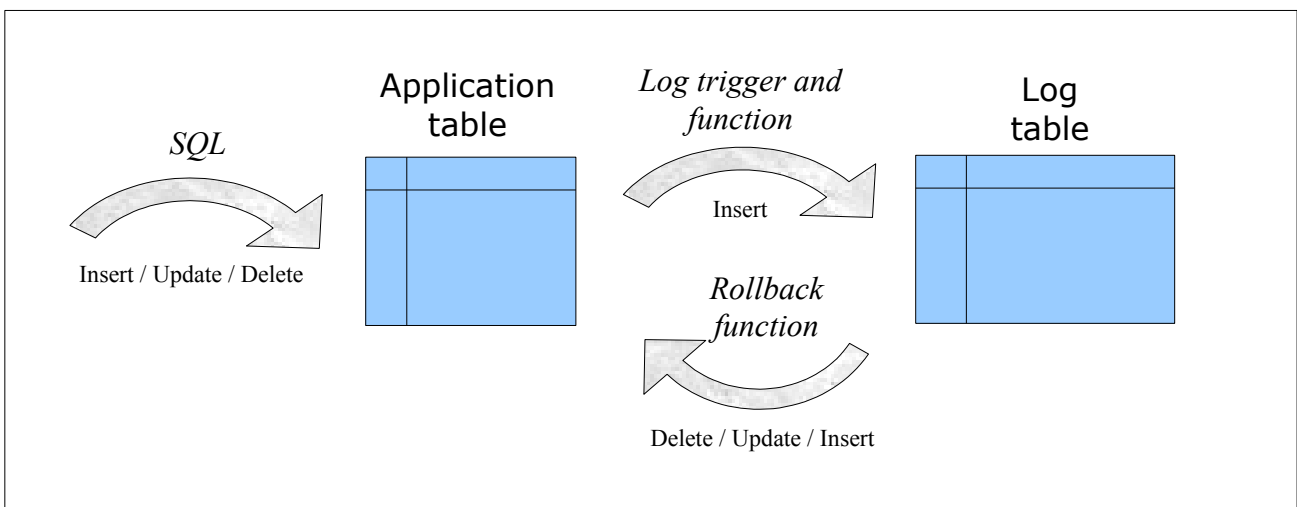
In order to be able to perform a rollback operation without having previously kept a physical image of the PostgreSQL cluster's files, all updates applied on application tables must be recorded, so that they can be cancelled.

With E-Maj, this updates recording takes the following form.

### 2.2.1 Created objects

For each application table, the following objects are created:

- a dedicated log table, containing data corresponding to the updates applied on the application table,
- a trigger and a specific function, that, for each row creation (*INSERT*), change (*UPDATE*) or suppression (*DELETE*), record into the log table all data needed to potentially cancel later this elementary action,
- a rollback function, that can cancel all or some recorded updates for the table, (except for tables belonging to group created in “*audit\_only*” mode),
- starting from PostgreSQL 8.4, another trigger that blocks any execution of a *TRUNCATE* SQL verb while log triggers are activated.



A log table has the same structure as its corresponding application table. However, it contains some additional technical columns:

- a unique identifier, as an integer associated to a sequence, and used as a primary key,
- the precise date and time of the update,
- the type of the executed SQL operation: INS for INSERT, UPD for UPDATE et DEL for DELETE,
- an attribute taking either 'OLD' or 'NEW' value, allowing to distinguish old and new values of updated rows,



- the internal transaction identifier (PostgreSQL *ctid*) that performed the update,
- the connection role who performed the update,
- the ip address of the user who performed the update.

To let E-Maj work, some other technical objects are also created at extension installation time:

- 10 tables,
- 2 types,
- 62 technical functions, 30 of them being directly callable by users,
- 1 specific schema, named *emaj*, that contains all these relational objects as well as the log tables and functions,
- 1 tablespace, *tspemaj*, hosting all these technical and log tables,
- 2 roles acting as groups (NOLOGIN): *emaj\_adm* to manage E-Maj components, and *emaj\_viewer* to only look at E-Maj components.

Some technical tables, whose structure is interesting to know, are described in the coming chapters.

## 2.2.2 Norm for E-Maj objects naming

All objects associated to application tables have names built with the name of their related table and schema. More precisely, for an application table in a given schema:

- the name of the log table is:  
`<schema.name>_<table.name>_log`
- the name of the log function is:  
`<schema.name>_<table.name>_log_fnct`
- the name of the log trigger is:  
`<schema.name>_<table.name>_emaj_log_trg`
- the name of the trigger that blocks *TRUNCATE* verb is:  
`<schema.name>_<table.name>_emaj_trunc_trg`
- the name of the rollback function is:  
`<schema.name>_<table.name>_rlbk_fnct`
- the name of the sequence associated to the log table is:  
`<schema.name>_<table.name>_log_emaj_id_seq`

Other E-Maj functions name is also normalised:

- functions names that begins with 'emaj\_' are functions that are callable by users,
- functions names that begins with '\_' are internal functions that should not be called directly.

## 3 HOW TO INSTALL E-MAJ

---

In this chapter, we will describe how to download and install the E-Maj extension. A last chapter deals with uninstall.

### 3.1 EXTENSION DOWNLOAD AND DECOMPRESSION

---

#### 3.1.1 Download

E-Maj is available for download on two Internet sites:

- the pgFoundry.org Internet site, at the url: <http://pgfoundry.org/projects/emaj/>.
- PGXN, PostgreSQL Extension Network, at the url <http://pgxn.org>.

#### 3.1.2 Decompression

The extension is delivered as a single compressed file. To be usable, this file must be decompressed.

Under Windows, you can use your favourite decompression utility. Under Unix/Linux, a command like :

```
tar -xvzf emaj-<version>.tar.gz
```

can be used for *.tar.gz* file or

```
unzip e-maj-<version>.zip
```

for a *.zip* file.

A new *emaj-<version>* directory is now available, containing the following files tree:

- *sql/emaj.sql* psql script to install the E-Maj components
- *sql/emaj.control\_base* extension control file base file
- *sql/emaj--0.10.1.sql* script for the CREATE EXTENSION statement
- *sql/emaj-0.10.0-to-0.10.1.sql* psql script to migrate E-Maj from version 0.10.0 to 0.10.1
- *sql/emaj--unpackaged--0.10.1.sql* script to transform installed E-Maj into extension
- *sql/emaj--0.10.0--0.10.1.sql* script to migrate E-Maj extension from version 0.10.0 to 0.10.1
- *sql/demo.sql* psql E-Maj demonstration script
- *sql/prep-pr.sql* psql test script for parallel rollbacks
- *sql/uninstall.sql* psql script to uninstall the E-Maj components

- README reduced extension's documentation
- CHANGES release notes
- LICENSE information about E-Maj license
- AUTHORS who are the authors
- META.json technical data for PGXN
- doc/Emaj.<version>\_doc\_en.pdf English version of the full E-Maj documentation
- doc/Emaj.<version>\_doc\_fr.pdf French version of the full E-Maj documentation
- doc/Emaj.<version>.en.pdf presentation of the E-Maj extension
- php/emajParallelRollback.php php tool for parallel rollback

## 3.2 E-MAJ EXTENSION SETUP

---

If E-Maj is already installed in the database, please go on with §3.2.2.

### 3.2.1 Full installation

For the main installation step, it is necessary to distinguish PostgreSQL versions, version 9.1 and later taking benefit from the integrated *extension* architecture. But some preliminary operations are required.

#### 3.2.1.1 Preliminary operations

For these operations, the user must log on the concerned database as a superuser, with *psql* for instance.

If the PL/PGSQL language is not activated (it is not activated by default with PostgreSQL versions prior 9.0), it must be activated by the following SQL command:

```
CREATE LANGUAGE plpgsql;
```

The second preliminary operation consists in creating the tablespace named *tspemaj* that will be dedicated to E-Maj, except if this tablespace has already been created by a previous installation.

To achieve this, the associated storage space (a directory for Unix/Linux, or a folder for Windows) must first be created, this storage space being left empty. Then the following SQL command must be executed:

```
CREATE TABLESPACE tspemaj LOCATION '<tablespace.directory/folder>;'
```

For performance reasons, it is recommended to put the *tspemaj* tablespace and the application tables on separate disk spaces.

### 3.2.1.2 E-Maj components installation with PostgreSQL versions 8.2 to 9.0

The E-Maj components can now be installed into the database, by executing under psql the supplied *emaj.sql* script.

```
\i <emaj_directory>/sql/emaj.sql
```

To start with, the script verifies that the PostgreSQL version is at least 8.2, that the current user has the super-user attribute and that *tspemaj* tablespace already exists.

Then the script creates the *emaj* schema with its technical tables, types and functions.

If they are not already present, both *emaj\_adm* and *emaj\_viewer* roles are also created.

Finally, the installation script looks at the cluster configuration and may display a warning message regarding the *-max-prepared-statements* parameter (see §4.7.2).

At the end of its execution, the script displays the following message:

```
>>> E-Maj objects successfully created
```

### 3.2.1.3 E-Maj components installation with PostgreSQL versions 9.1 and later

Starting from PostgreSQL 9.1, it is possible to install E-Maj as an *extension*, within the meaning of PostgreSQL.

As a first step, some parameters related to the E-Maj extension must be supplied to PostgreSQL. Concretely, a file named:

```
share/postgresql/extension/emaj.control
```

must be stored into the PostgreSQL installation directory, based on the *sql/emaj.control\_base* file supplied within E-Maj.

In this file, only the *directory* parameter must be changed, by specifying the directory pathname that hosts installation scripts (typically the */sql* directory created at E-Maj decompression). Other parameters must be kept unchanged.

Then just create the *extension* by issuing the following statement:

```
CREATE EXTENSION emaj;
```

Note that the E-Maj extension is always installed into the *emaj* schema. It is not possible to specify another schema name in the *CREATE EXTENSION* statement.

The extension creation performs the same checks and creates the same objects (tables, types, functions and roles) than *emaj.sql* script.

If E-Maj 0.10.0 has been installed with *emaj.sql* script, it is possible then to transform the installation into an *extension*. To do this, just execute:

```
CREATE EXTENSION emaj FROM unpackaged;
```

Once installed as an *extension*, E-Maj can be uninstalled using a *DROP EXTENSION* statement.

### 3.2.2 Update an existing E-Maj version prior 0.10.0

Due to the large differences between version 0.10.1 and the versions prior 0.10.0, it is not interesting to get a simple migration path. For this reason, migrate from an E-Maj version prior 0.10 needs a deletion and re-installation of the extension. However, as E-Maj tablespace and roles can remain, it is not necessary to perform a full uninstall.

#### 3.2.2.1 Save user data

It may be useful to save the content of the *emaj\_group\_def* table in order to be able to easily reload it after the version update, by copying it outside the cluster with a *lcopy* command, or by duplicating the table outside the *emaj* schema with a SQL statement like:

```
CREATE TABLE public.savegroupdef AS SELECT * FROM emaj.emaj_group_def;
```

The same way, if the E-Maj administrator had changed parameters value into *emaj\_param* table, it may also be useful to keep a trace of these changes, for instance with:

```
CREATE TABLE public.saveparam AS SELECT * FROM emaj.emaj_param;
```

### 3.2.2.2 Reinstall E-Maj

With PostgreSQL version 8.2 to 9.0, just execute *emaj.sql* script as indicated in §3.2.1.2 to install the new 0.10.0 E-Maj version, the script deleting the previous version.

With 9.1 and later PostgreSQL versions, the *emaj* schema must be dropped first:

```
DROP SCHEMA emaj CASCADE;
```

This dropped the *extension* as well. Then create the *extension* as indicated in §3.2.1.3.

### 3.2.2.3 Restore user data

Data previously saved can now be restored into E-Maj tables.

The statement inserting data into *emaj\_group* table must take into account that the new table contains an additional column, that may be set to *NULL* (see §4.2.2).



Be careful also that in the new *emaj\_param* table, parameters that keep their default values have no corresponding rows any more. So insert into this table nothing but parameters whose value are customized to your installation (see §5.1).

## 3.2.3 Migration from an already installed 0.10.0 version

If a 0.10.0 E-Maj version is already installed, it is possible to perform a simple E-Maj update.

This update can be done without dropping existing tables groups, and even without stopping them if they are in active state. This means that updates on application tables can continue to be recorded during and after this version change. In such a case, a « *rollback* » on a mark set before the version change can be performed too.

Nevertheless, already created tables groups will not take benefit of all enhancements brought by this version until they will be dropped and re-created. This concerns in particular:

- the addition of a primary key on each log table to speed up some operations,
- the fix of some issues encountered with schema, table or column names containing some unusual characters like simple or double quotes.

### 3.2.3.1 Updating E-Maj components with PostgreSQL versions 8.2 to 9.0

To update E-Maj when using PostgreSQL version 8.2 to 9.0, the *emaj-0.10.0-to-0.10.1.sql* delivered psql script must be executed.

```
\i <emaj_directory>/sql/emaj-0.10.0-to-0.10.1.sql
```

The script reports the list of tables groups that will need to be recreated after the installation to take benefit of all enhancements brought by the new version.

At the end of its execution, the script displays the following message :

```
>>> E-Maj successfully migrated to 0.10.1
```

### 3.2.3.2 Updating E-Maj components with PostgreSQL versions 9.1+

As a first step, the:

```
share/postgresql/extension/emaj.control
```

file located in the PostgreSQL installation directory must be modified, to define the new *directory* that contains E-Maj components. The delivered *sql/emaj.control\_base* file can be used as template.

Then, the *extension* version can be changed using the following statement:

```
ALTER EXTENSION emaj UPDATE TO '0.10.1';
```

### 3.2.4 Changes in postgresql.conf configuration file

Main E-Maj functions set a lock on each table of a processed tables group. If some groups contains a large number of tables, it may be necessary to increase the value of the *max\_locks\_per\_transaction* parameter in the *postgresql.conf* configuration file. The default value of this parameter equals 64. It can be increased up to a value at least equal to the number of tables of the largest group.

Furthermore, if the parallel rollback tool may be used (see § 4.7), it will be probably necessary to adjust the *max\_prepared\_transaction* parameter.

### 3.2.5 Test and demonstration

It is possible to check whether the installed E-Maj components work fine, by executing under *psql* the *demo.sql* script.

```
\i <emaj_directory>/sql/demo.sql
```

If no error is encountered, the final message

```
"--- End of E-Maj demo ---"
```

is displayed.

### 3.3 E-MAJ UNINSTALL

---

For this operation, the user must log on the concerned database with `psql`, as a superuser.

If some tables groups in logging state remain, they must be stopped with the `emaj_stop_group()` function (see § 4.2.8).

If the drop of the `emaj_adm` and `emaj_viewer` roles is desirable, rights on them given to other roles must be previously deleted, using `REVOKE` SQL verbs.

```
REVOKE emaj_adm FROM <role.or.roles.list>;  
REVOKE emaj_viewer FROM <role.or.roles.list>;
```

If these `emaj_adm` and `emaj_viewer` roles own access rights on other application objects, these rights must be suppressed too, before starting the uninstall operation.

Then, the `uninstall.sql` script delivered with the installed E-Maj version has to be executed.

```
\i <emaj_directory>/uninstall.sql
```

This script first drops the `emaj` schema with all contained objects. If E-Maj has been installed as an `EXTENSION`, this PostgreSQL object is also dropped.

If `emaj_adm` and `emaj_viewer` roles are not associated to other role any longer, and do not own rights on other tables, they are dropped. Otherwise a message is displayed, inviting to a manual deletion.

The `tspemaj` tablespace is not dropped by the script. To drop it, the only SQL command to execute is:



```
DROP TABLESPACE tspemaj;
```



If E-Maj is installed into several databases of the same PostgreSQL cluster, the *tspemaj* tablespace is used by all E-Maj objects of all databases. So it could only be dropped after the deletion of all E-Maj components of all the cluster's databases.

## 4 HOW TO USE E-MAJ

---

### 4.1 SET-UP THE E-MAJ ACCESS POLICY

---

A bad usage of E-Maj can break databases integrity. So it is advisable to only authorise its use to specific skilled users.

#### 4.1.1 E-Maj roles

To use E-Maj, it is possible to log on as superuser. But for safety reason, it is preferable to take advantage of both roles created by the installation script:

- *emaj\_adm* is used as administration role ; it can execute all functions and access to all E-Maj tables, with reading and writing rights,
- *emaj\_viewer* is used for read only purpose ; it can only execute statistics functions and can only read E-Maj tables.

All rights given to *emaj\_viewer* are also given to *emaj\_adm*.

When created, these roles have no connection capability (no defined password and *NOLOGIN* option). It is recommended NOT to give them any connection capability. Instead, it is sufficient to give the rights they own to other roles, with *GRANT SQL* verbs.

#### 4.1.2 Giving E-Maj rights

Once logged on as superuser in order to have the sufficient rights, execute one of the following commands to give a role all rights associated to one of both *emaj\_adm* or *emaj\_viewer* roles:

```
GRANT emaj_adm TO <my.emaj.administrator.role>;  
GRANT emaj_viewer TO <my.emaj.viewer.role>;
```

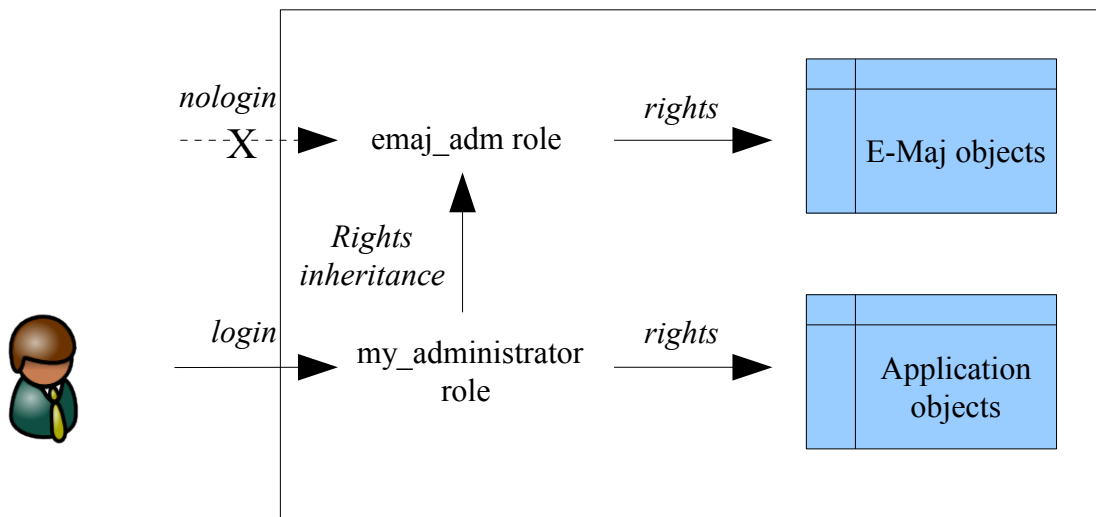
Of course, *emaj\_adm* or *emaj\_viewer* rights can be given to several roles.

### 4.1.3 Giving rights on application tables and objects

To let an E-Maj administrator also access to application tables or to other application objects (schemas, sequences, views, functions,...), it is possible to give rights on these objects to *emaj\_adm* or *emaj\_viewer* roles. But it is preferable to only give these rights to the roles which are also given *emaj\_adm* or *emaj\_viewer* rights, so that the E-Maj roles only directly own rights on E-Maj tables and objects.

### 4.1.4 Synthesis

The following schema represents the recommended rights organisation for an E-Maj administrator.



Of course the schema also applies to *emaj\_viewer* role.


Except when explicitly noticed, the operations presented later can be indifferently executed by a superuser or by a role belonging to the *emaj\_adm* group.

## 4.2 MAIN FUNCTIONS

Before describing each main E-Maj function, it is interesting to have a global view on the typical operations chain.

### 4.2.1 Operations chain

The possible chaining of operations for a tables group can be materialised by this schema.



## 4.2.2 Define tables groups

The content of tables groups E-Maj will manage has to be defined by populating the *emaj.emaj\_group\_def* table. One row has to be inserted into this table for each application table or sequence to include into a tables group. This *emaj.emaj\_group\_def* table contains 3 columns of type TEXT:

- *grpdef\_group* tables group name
- *grpdef\_schema* name of the schema containing the application table or sequence
- *grpdef\_tblseq* application table or sequence name
- *grpdef\_priority* priority level for the table or sequence in E-Maj processings

The administrator can populate this table by any usual mean: *INSERT* SQL verb, *COPY* SQL verb, *copy* psql command, graphic tool, etc.

A tables group name contains at least 1 character. It may contain spaces and/or any punctuation characters. But it is advisable to avoid commas, simple or double quotes.

A table or a sequence of a given schema cannot be affected to several tables groups. All tables of a schema are not necessarily member of the same group. Some of them can belong to another group. Some others can belong to any group.

All tables affected to a group not created in “*audit\_only*” mode must have an explicit primary key (*PRIMARY KEY* clause in *CREATE TABLE* or *ALTER TABLE*).

If a sequence is associated to an application table, it must be explicitly declared as member of the same group as its table, so that, in case of rollback, the sequence can be reset to its state at the set mark time.

On the contrary, log tables and their sequence should NOT be referenced in a tables group!

The type of the *grpdef\_priority* column is *INTEGER* and may be NULL. It defines a priority order in E-Maj tables processings. This can be useful at table lock time. Indeed, by locking tables in the same order as what is typically done by applications, it may reduce the risk of deadlock. E-Maj functions process tables in *grpdef\_priority* ascending order, NULL being processed last. For a same priority level, tables are processed in alphabetic order of schema name and table name.

The content of the *emaj\_group\_def* table is case sensitive. Schema names, table names and sequence names must reflect the way PostgreSQL registers them in its catalogue. These names are mostly in lower case. But if a name is encapsulated by double quotes in SQL statements because it contains any upper case characters or spaces, then it must be registered into the *emaj\_group\_def* table with the same upper case characters or spaces.



To guarantee the integrity of tables managed by E-Maj, it is essential to take a particular attention to this tables groups content definition step. If a table were missing, its content would be out of synchronisation with other tables it is related to, after a rollback operation. In particular, when application tables are created or suppressed, it is important to always maintain an up-to-date content of this *emaj\_group\_def* table.

### 4.2.3 Create a tables group

Once the content of a tables group is defined, E-Maj can create the group. To do this, there is only one SQL statement to execute:

```
SELECT emaj.emaj_create_group('<group.name>',<is_rollbackable>);
```

or in an abbreviate form:

```
SELECT emaj.emaj_create_group('<group.name>');
```

The second parameter, boolean, indicates whether the group is a “*rollbackable*” (with value true) or an “*audit\_only*” (with value false) group. If this second parameter is not supplied, the group is considered “*rollbackable*”.

The function returns the number of tables and sequences contained by the group.

For each table of the group, this function creates the associated log table, the log function and trigger, as well as the trigger that blocks the execution of *TRUNCATE* SQL statements (starting PostgreSQL 8.4). If the group is “*rollbackable*”, the rollback function is also created.

The *emaj\_create\_group()* function also checks the existence of application triggers on any tables of the group. If a trigger exists on a table of the group, a message is returned, suggesting the user to verify that this trigger does not update any tables that would not belong to the group.

All actions that are chained by the *emaj\_create\_group()* function are executed on behalf of a unique transaction. As a consequence, if an error occurs during the operation, all tables, functions and triggers already created by the function are cancelled.

By registering the group composition in the *emaj\_relation* internal table, the *emaj\_create\_group()* function freezes its definition for the other E-Maj functions, even if the content of the *emaj\_group\_def* table is modified later.

A tables group can be suppressed by the *emaj\_drop\_group()* function (see § 4.2.9).

### 4.2.4 Start a tables group

Starting a tables group consists in activating the recording of updates for all tables of the group. To achieve this, the following command must be executed:

```
SELECT emaj.emaj_start_group('<group.name>', '<mark.name>');
```

The group must be first in *IDLE* state.

A mark name must be specified. It will be the first mark on which a rollback will be later possible.

The mark name may contain a generic '%' character. Then this character is replaced by the current transaction start time, with the pattern « hh.mn.ss.mmm »,

If the parameter representing the mark is empty or *NULL*, a name is automatically generated: « *MARK\_%* », where the '%' character represents the current transaction start time.

The function returns the number of tables and sequences contained by the group.

To be sure that no transaction implying any table of the group is currently running, the *emaj\_start\_group()* function explicitly set an *EXCLUSIVE* lock on each table of the group. If transactions accessing these tables are running, this can lead to *deadlock*. If the *deadlock* processing impacts the execution of the E-Maj function, the error is trapped and the lock operation is repeated, with a maximum of 5 attempts.

Before activating log triggers that will log updates into log tables, all log tables are emptied.

When a group is started, its state becomes « *LOGGING* ».

#### 4.2.5 Set an intermediate mark

When all tables and sequences of a group are considered as being in a stable state that can be used for a potential rollback, a mark can be set. This is done with the following SQL statement:

```
SELECT emaj.emaj_set_mark_group('<group.name>', '<mark.name>');
```

The tables group must be in *LOGGING* state.

A mark having the same name can not already exist for this tables group.

The mark name may contain a generic '%' character. Then this character is replaced by the current transaction start time, with the pattern « hh.mn.ss.mmm »,

If the parameter representing the mark is empty or *NULL*, a name is automatically generated: « *MARK\_%* », where the '%' character represents the current transaction start time.

The function returns the number of tables and sequences contained by the group.

The *emaj\_set\_mark\_group()* function records the identity of the new mark, with the state of the application sequences belonging to the group, as well as the state of the log sequences associated to each table of the group.

Note that it is possible to set two consecutive marks without any update on any table between these marks.

The *emaj\_set\_mark\_group()* function set *ROW EXCLUSIVE* locks on each table of the group in order to be sure that no transaction having already performed updates on any table of the group is running. However, this does not guarantee that a transaction having already read one or several tables before the mark set, updates tables after the mark set. In such a case, these updates would be candidate for a potential rollback to this mark.

#### 4.2.6 Rollback a tables group

If it is necessary to reset tables and sequences of a group in the state they were when a mark was set, a rollback must be performed. To perform a simple (“*unlogged*”) rollback, the following SQL statement can be executed:

```
SELECT emaj.emaj_rollback_group('<group.name>', '<mark.name>');
```

The tables group must be in *LOGGING* state.

The '*EMAJ\_LAST\_MARK*' keyword can be used as mark name, meaning the last set mark.

The function returns the number of tables and sequences that have been **effectively** modified by the rollback operation.

To be sure that no transaction implying any table of the group is currently running, and that no other transaction will be able to access any table during the rollback operation, the *emaj\_rollback\_group()* function explicitly set an *EXCLUSIVE* lock on each table of the group. If transactions accessing these tables are running, this can lead to *deadlock*. If the *deadlock* processing impacts the execution of the E-Maj function, the error is trapped and the lock operation is repeated, with a maximum of 5 attempts.

If tables belonging to the group to rollback have triggers, it may be necessary to deactivate them before the rollback and re-activate them after (see §5.7.3).



When the rollback operation is completed, are deleted:

- all log tables rows corresponding to the rolledback updates,
- all marks later than the mark referenced in the rollback operation.

Then, it is possible to continue updating processings, to set other marks, and if needed to perform another rollback at any mark.



By nature, the reset of sequence is not “cancellable” in case of abort and rollback of the transaction that executes the *emaj\_rollback\_group()* function. That is the reason why the processing of application sequences is always performed after the processing of application tables. However, even-though the time needed to rollback a sequence is very short, a problem may occur during this last phase. Rerunning immediately the *emaj\_rollback\_group()* function would not break database integrity. But any other database access before the second execution may lead to wrong values for some sequences.

## 4.2.7 Perform a logged rollback of a tables group

Another function executes a “logged” rollback. In this case, log triggers on application tables are not disabled during the rollback operation. As a consequence, the updates on application tables are also recorded into log tables, so that it is possible to cancel a rollback. In other words, it is possible to rollback ... a rollback.

To execute a “logged” rollback, the following SQL statement can be executed:

```
SELECT emaj.emaj_logged_rollback_group('<group.name>', '<mark.name>');
```

The usage rules are the same as with *emaj\_rollback\_group()* function.

The tables group must be in *LOGGING* state.

The 'EMAJ\_LAST\_MARK' keyword can be used as mark name, meaning the last set mark.

The function returns the number of tables and sequences that have been **effectively** modified by the rollback operation.

To be sure that no transaction implying any table of the group is currently running, and that no other transaction will be able to access any table during the rollback operation, the *emaj\_rollback\_group()* function explicitly set an *EXCLUSIVE* lock on each table of the group. If transactions accessing these tables are running, this can lead to *deadlock*. If the *deadlock* processing impacts the execution of the E-Maj function, the error is trapped and the lock operation is repeated, with a maximum of 5 attempts.

If tables belonging to the group to rollback have triggers, it may be necessary to deactivate them before the rollback and re-activate them after (see §5.7.3).

Unlike with *emaj\_rollback\_group()* function, at the end of the operation, the log tables content as well as the marks post the rollback mark remain.

At the beginning and at the end of the operation, the function automatically sets on the group two marks named:

- 'RLBK\_<rollback.mark>\_<rollback.time>\_START'
- 'RLBK\_<rollback.mark>\_<rollback.time>\_DONE'

where *rollback.time* represents the start time of the transaction performing the rollback, expressed as "hours.minutes.seconds.milliseconds".

Then, it is possible to continue updating processings, to set other marks, and if needed to perform another rollback at any mark, including the mark set at the beginning of the rollback, to cancel it, or even an old mark that was set after the mark used for the rollback.

Rollback from different types (*logged/unlogged*) may be executed in sequence. For instance, it is possible to chain the following steps:

```
Set Mark M1
...
Set Mark M2
...
Logged Rollback to M1,
generating RLBK_M1_<time>_STRT,
and RLBK_M1_<time>_DONE
...
Rollback to RLBK_M1_<time>_DONE
(to cancel the updates performed after the first rollback)
...
Rollback to RLBK_M1_<time>_STRT
(to finally cancel the first rollback)
```

## 4.2.8 Stop a tables group

When one wishes to stop the updates recording for tables of a group, it is possible to deactivate the logging mechanism, using the command:

```
SELECT emaj.emaj_stop_group('<group.name>');
```

The function returns the number of tables and sequences contained by the group.

Stopping a tables group simply deactivate log triggers of application tables of the group. The locks set can lead to *deadlock*. If the *deadlock* processing impacts the execution of the E-Maj function, the error is trapped and the lock operation is repeated, with a maximum of 5 attempts.

Additionally, the *emaj\_stop\_group()* function changes the status of all marks set for the group into a *DELETED* state. Then, it is not possible to execute a rollback command any more, even though no updates have been applied on tables between the execution of both *emaj\_stop\_group()* and *emaj\_rollback\_group()* functions.

But the content of log tables and E-Maj technical tables can be examined.

When a group is stopped, its state becomes « *IDLE* » again.

Executing twice the *emaj\_stop\_group()* function for a tables group already stopped does not generate an error. Only a warning message is returned.

## 4.2.9 Drop a tables group

To drop a tables group previously created by the *emaj\_create\_group()* function, this group must be already in idle state. If it is not the case, the *emaj\_stop\_group()* function has to be used (see § 4.2.8).

Then, just execute the SQL command:

```
SELECT emaj.emaj_drop_group('<group.name>');
```

The function returns the number of tables and sequences contained by the group.

For this tables group, the *emaj\_drop\_group()* function drops all the objects that have been created by the *emaj\_create\_group()* function: log tables, log and rollback functions, log triggers.

The locks set by this operation can lead to *deadlock*. If the *deadlock* processing impacts the execution of the E-Maj function, the error is trapped and the lock operation is repeated, with a maximum of 5 attempts.

## 4.2.10 Change a tables group content

The content of a tables group may vary. It may be necessary to add new tables or sequences into a group, or suppress some others from a group. In that case, the following operations must be chained:

- stop the group if it is in logging state, using the *emaj\_stop\_group()* function,
- drop the group with the *emaj\_drop\_group()* function,
- adapt the content of the *emaj\_group\_def* table to reflect the desired change in the group's content,
- recreate the group with the *emaj\_create\_group()* function.

It is possible to anticipate the update of the *emaj\_group\_def* table, even if the tables group is in logging state.

#### **4.2.11 Change the structure of an application table**

If the structure of an application table belonging to a tables group changes (addition or suppression of a column, change in a column type), it is mandatory to drop the related tables group, using the *emaj\_drop\_group()* function, and then to recreate it, using the *emaj\_create\_group()* function. Indeed, these evolutions will have an impact on the associated log table structure.

In case of discrepancy between the structure of both application and related log tables, E-Maj generates an error at start group time, or set mark time or rollback time.

## 4.3 MULTI-GROUPS FUNCTIONS

---

### 4.3.1 General information

To be able to synchronize current operations like group start or stop, set mark or rollback, usual functions dedicated to these tasks have twin-functions that process several tables groups in a single call.

The resulting advantages are:

- to process all tables group in a single transaction,
- to lock tables belonging to all groups at the beginning of the operation, to minimize the risk of deadlock.

### 4.3.2 Functions list

There are five functions processing several groups:

- *emaj.emaj\_start\_groups(<groups.array>,<start.mark>)* starts several groups,
- *emaj.emaj\_stop\_groups(<groups.array>)* stops several groups,
- *emaj.emaj\_set\_mark\_groups(<groups.array>,<mark>)*, sets a mark on several groups,
- *emaj.emaj\_rollback\_groups(<groups.array>,<mark>)* process a simple rollback for several groups,
- *emaj.emaj\_logged\_rollback\_groups(<groups.array>,<mark>)* process a “logged” rollback for several groups.

### 4.3.3 Syntax for groups array

The SQL type of the <groups.array> parameter passed to the multi-groups functions is *TEXT[ ]*, i.e. an array of text data.

According to SQL standard, there are 2 possible syntaxes to specify a groups array, using either braces { }, or the *ARRAY* function.

When using { and }, the full list is written between simple quotes, then braces frame the comma separated elements list, each element been placed between double quotes. For instance, in our case, we can write :

```
' { "group 1" , "group 2" , "group 3" } '
```

The SQL function *ARRAY* builds an array of data. The list of values is placed between brackets [ ], and values are separated by comma. For instance, in our case, we can write :

```
ARRAY [ 'group 1' , 'group 2' , 'group 3' ]
```

Both syntax are equivalent.

#### 4.3.4 Other considerations

The order of the groups in the groups list is not meaningful. During the E-Maj operation, the processing order of tables only depends on the priority level defined for each table, and, for tables having the same priority level, from the alphabetic order of their schema and table names.

It is possible to call a multi-groups function to process a list of ... one group, or even an empty list. So a set oriented build of this list is possible (using for instance the *array\_agg()* function that is available from PostgreSQL version 8.4).

It is also permitted to have a list with duplicate values, NULL values or empty strings. In all these cases, the value is ignored and a warning message is generated, like in case of empty list.

Format and usage of these functions are strictly equivalent to those of their twin-functions.

However, an additional condition exists for rollback functions. The supplied mark must correspond to the same point in time for all groups. In other words, this mark must have been set by the same *emaj\_set\_mark\_group()* function call.

## 4.4 MARKS MANAGEMENT FUNCTIONS

---

### 4.4.1 Comments on marks

In order to set a comment on any mark, the following statement can be executed:

```
SELECT emaj.emaj_comment_mark_group('<group.name>', '<mark>',  
'<comment>');
```

The function doesn't return any data.

To modify an existing comment, just call the function again for the same tables group and the same mark, with the new comment.

To delete a comment, just call the function, supplying a NULL value as comment.

Comments are stored into the *mark\_comment* column from the *emaj\_mark* table, which describes ... marks.

Comments are mostly interesting when using the E-Maj phpPgAdmin plug-in (See §6). Indeed, the plug-in systematically displays the comments in the groups marks list.

### 4.4.2 Rename a mark

A mark that has been previously set by one of both *emaj\_create\_group()* or *emaj\_set\_mark\_group()* functions can be renamed, using the SQL statement:

```
SELECT emaj.emaj_rename_mark_group('<group.name>', '<mark.name>',  
'<new.mark.name>');
```

The keyword *'EMAJ\_LAST\_MARK'* can be used as mark name. It then represents the last set mark.

The function does not return any data.

A mark having the same name as the requested new name should not already exist for the tables group.

### 4.4.3 Delete a mark

A mark can also be deleted, using the SQL statement:

```
SELECT emaj.emaj_delete_mark_group('<group.name>', '<mark.name>');
```

The keyword `'EMAJ_LAST_MARK'` can be used as mark name. It then represents the last set mark.

The function returns 1, corresponding to the number of effectively deleted marks.

As at least one mark must remain after the function has been performed, a mark deletion is only possible when there are at least two marks for the concerned tables group.

If the deleted mark is the first mark of the tables group, the useless rows of log tables are deleted.

### 4.4.4 Delete oldest marks

To easily delete in a single operation all marks prior a given mark, the following statement can be executed:

```
SELECT emaj.emaj_delete_before_mark_group('<group.name>',  
'<mark.name>');
```

The function deletes all marks prior the supplied mark, this mark becoming the new first available mark. It also suppresses from log tables all rows related to the deleted period of time.

With this function, it is quite easy to use E-Maj during a long period of time, without stopping and restarting groups, while limiting the disk space needed for `tspemaj` tablespace.

However, as the log rows deletion cannot use any `TRUNCATE` command (as with the `emaj_start_group()` or `emaj_reset_group()` functions), using `emaj_delete_before_group()` function may take a longer time than simply stopping and restarting the group. In return, no lock are set on the tables of the group.



## 4.5 STATISTICS FUNCTIONS

---

There are two functions that return statistics on log tables content:

- *emaj\_log\_stat\_group()* quickly delivers, for each table of a group, the number of updates that have been recorded in the related log tables, either between 2 marks or since a particular mark,
- *emaj\_detailed\_log\_stat\_group()* brings a more detailed information than *emaj\_log\_stat\_group()*, the number of updates been reported per table, SQL type (INSERT/UPDATE/DELETE) and connection role.

Another E-Maj function, *emaj\_estimate\_rollback\_duration()*, provides an estimate of how long a rollback for a group to a given mark may last.

Finally, a function named *emaj\_get\_previous\_mark\_group()*, returns for a group the name of the latest mark preceding a given date and time.

Both functions can be used by *emaj\_adm* and *emaj\_viewer* E-Maj roles.

### 4.5.1 Global statistics about logs

Full global statistics about logs content are available with this SQL statement:

```
SELECT * FROM emaj.emaj_log_stat_group('<group.name>', '<start.mark or NULL>', '<end.mark or NULL>');
```

The function returns a set of rows, whose type is named *emaj.emaj\_log\_stat\_type*, and contains the following columns:

- *stat\_group* : tables group name (type TEXT),
- *stat\_schema* : schema name (type TEXT),
- *stat\_table* : table name (type TEXT),
- *stat\_rows* : number of updates recorded into the related log table (type BIGINT)

A NULL value or an empty string (") supplied as start mark represents the oldest mark.

A NULL value supplied as end mark represents the current situation.

The keyword '*EMAJ\_LAST\_MARK*' can be used as mark name. It then represents the last set mark.

The function returns one row per table, even if there is no logged update for this table. In this case, *stat\_rows* columns value is 0.

It is possible to easily execute more precise requests on these statistics. For instance, once the *test-emaj-2.sql* test script has been executed, it is possible to get the number of database updates by application schema, with a statement like:

```
postgres=# SELECT stat_schema, sum(stat_rows)
FROM emaj.emaj_log_stat_group('myAppl1', NULL, NULL)
GROUP BY stat_schema;
 stat_schema | sum
-----+-----
myschema    |  41
(1 row)
```

There is no need for log table scans to get these statistics. For this reason, they are delivered quickly.

But returned values may be approximative (in fact over-estimated). This occurs in particular when transactions executed between both requested marks have performed tables updates before been cancelled.

#### 4.5.2 Detailed statistics about logs

Scanning log tables brings a more detailed information, at a higher response time cost. So can we get fully detailed statistics with the following SQL statement:

```
SELECT * FROM emaj.emaj_detailed_log_stat_group('<group.name>',
'<start.mark or NULL>', '<end.mark or NULL>');
```

The function returns a set of rows, whose type is named *emaj.emaj\_detailed\_log\_stat\_type*, and contains the following columns:

- *stat\_group* : tables group name (type TEXT),
- *stat\_schema* : schema name (type TEXT),
- *stat\_table* : table name (type TEXT),
- *stat\_role* : connection role (type VARCHAR(32)),
- *stat\_verb* : type of the SQL verb that has performed the update (type VARCHAR(6), with values: *INSERT* / *UPDATE* / *DELETE*),
- *stat\_rows* : number of updates recorded into the related log table (type BIGINT)

A NULL value or an empty string (") supplied as start mark represents the oldest mark.

A NULL value supplied as end mark represents the current situation.

The keyword *'EMAJ\_LAST\_MARK'* can be used as mark name. It then represents the last set mark.

Unlike *emaj\_log\_stat\_group*, *emaj\_detailed\_log\_stat\_group* function doesn't return any row for tables having no logged updates inside the requested marks range. So *stat\_rows* column never contains 0.

It is possible to easily execute more precise requests on these statistics. For instance, once the *test-emaj-2.sql* test script has been executed, it is possible to get the number of updates for a given table, here *mytbl1*, per SQL verb, using a statement like:

```
postgres=# SELECT stat_table, stat_verb, stat_rows
FROM emaj.emaj_detailed_log_stat_group('myApp11', NULL, NULL)
WHERE stat_table='mytbl1';
 stat_table | stat_verb | stat_rows
-----+-----+-----
 mytbl1    | DELETE   |         1
 mytbl1    | INSERT   |         6
 mytbl1    | UPDATE   |         2
(3 rows)
```

### 4.5.3 Estimate the rollback duration

The *emaj\_estimate\_rollback\_duration()* function returns an idea of the time needed to rollback a group to a given mark. It can be called with a statement like:

```
SELECT emaj.emaj_estimate_rollback_duration('<group.name>',
'<mark.name>');
```

The keyword *'EMAJ\_LAST\_MARK'* can be used as mark name. It then represents the last set mark.

The function returns an *INTERVAL* data.

The tables group must be in *IDLE* state.

This duration estimate is approximative. It takes into account:

- the number of updates in log tables to process, as returned by the *emaj\_log\_stat\_group()* function,
- recorded duration of already performed rollbacks for the same tables,
- 4 generic parameters (see § 5.1) that are used as default values when no statistics have been already recorded for the tables to process.

The precision of the result cannot be high. The first reason is that, INSERT, UPDATE and DELETE having not the same cost, the part of each SQL type may vary. The second reason is that the load of the server at rollback time can be very different from one run to another. However, if there is a time constraint, the order of magnitude delivered by the function can be helpful to determine of the rollback operation can be performed in the available time interval.

If no statistics on previous rollbacks are available and if the results quality is poor, it is possible to adjust parameters listed in chapter 5.1. It is also possible to manually change the *emaj.emaj\_rlbk\_stat* table's content that keep a trace of the previous rollback durations, for instance by deleting rows corresponding to rollback operations performed in unusual load conditions.

#### 4.5.4 Get a mark name

The *emaj\_get\_previous\_mark\_group()* function provides the name of the latest mark before a given date and time for a tables group.

```
SELECT emaj.emaj_get_previous_mark_group('<group.name>', '<date.time>');
```

The date and time must be expressed as a *TIMESTAMPTZ* datum, for instance the literal '2011/06/30 12:00:00 +02'.

If the supplied time strictly equals the time of an existing mark, the returned mark would be the preceding one.

This function is particularly useful when used with the *emaj\_delete\_before\_mark\_group()*. For example, to suppress all marks (and the associated log rows) set since more than 24 hours, the following statement can be executed:

```
SELECT emaj.emaj_delete_before_mark_group('<group>',  
emaj.emaj_get_previous_mark_group('<group>', current_timestamp - '1  
DAY'::INTERVAL));
```

## 4.6 OTHER FUNCTIONS

---

### 4.6.1 Simultaneous rollback and stop of a tables group

If the E-Maj administrator wants to rollback a table group and to stop this same group just after, it is possible to group actions performed by both *emaj\_rollback\_group()* and *emaj\_stop\_group()* functions into a single operation.

```
SELECT emaj.emaj_rollback_and_stop_group('<group.name>', '<mark.name>');
```

As with *emaj\_rollback\_group()* function, this function returns the number of tables and sequences that have been effectively modified by the rollback operation.

Using the *emaj\_rollback\_and\_stop\_group()* function instead of both usual functions saves time. Indeed, the group being stopped, no rollback will be possible after the completion of the function. As a consequence, the *emaj\_rollback\_and\_stop\_group()* function do not erase the content of log tables corresponding to the cancelled updates. The time saved will be all the more significant as the rollback cancels a large number of updates.

### 4.6.2 Reset log tables of a group

In standard use, all log tables of a tables group are purged at *emaj\_start\_group()* time. But, if needed, it is possible to reset log tables, using the following SQL statement:

```
SELECT emaj.emaj_reset_group('<group.name>');
```

The function returns the number of tables and sequences contained by the group.

Of course, in order to reset log tables, the tables group must be in *idle* state.

### 4.6.3 Comments on groups

In order to set a comment on any group, the following statement can be executed:

```
SELECT emaj.emaj_comment_group('<group.name>', '<comment>');
```

The function doesn't return any data.

To modify an existing comment, just call the function again for the same tables group, with the new comment.

To delete a comment, just call the function, supplying a NULL value as comment.

Comments are stored into the *group\_comment* column from the *emaj\_group* table, which describes ... groups.

#### 4.6.4 Check the consistency of the E-Maj environment

A function is also available to check the consistency of the E-Maj environment, including a global check of the environment (orphan tables and functions in the *emaj* schema) and an integrity check of all created tables groups. This function can be called with the following SQL statement:

```
SELECT * FROM emaj.emaj_verify_all();
```

The global check verifies that:

- each log table existing in the *emaj* schema corresponds to a table referenced into the *emaj\_group* table, the internal table that describes the composition of created groups,
- each log or rollback function that exists in the *emaj* schema corresponds to a table referenced in this same *emaj\_group* table.

Then, for each created tables group, the function performs the same checks than those performed when a group is started, a mark is set or a rollback is executed (see §5.2).

The function returns a set of rows describing the detected discrepancies. If no error has been detected, the function returns a set of rows containing the following messages:

*'Global checking: no error encountered'*

then, for each tables group:

*'Checking <tables.group.name>: no error encountered'*

The *emaj\_verify\_all()* function can be executed by any role belonging to *emaj\_adm* or *emaj\_viewer* roles.

If errors are detected, for instance after an application table referenced in a tables group has been dropped, appropriate measures must be taken. Typically, the potential orphan log tables or functions must be manually dropped.

## 4.6.5 Forced suppression of a tables group

It may happen that a damaged tables group cannot be stopped. But not being stopped, it cannot be dropped. To be able to drop a tables group while it is still in logging state, a special function exists.

```
SELECT emaj.emaj_force_drop_group('<group.name>');
```

The function returns the number of tables and sequences contained by the group.

This *emaj\_force\_drop\_group()* functions performs the same actions than the *emaj\_drop\_group()* function, but without checking the state of the group. So, it is recommended to only use this function if it is really needed.

## 4.6.6 Snap tables of a group

It may be useful to take images of all tables and sequences belonging to a group to be able to analyse their content or compare them. It is possible to dump on files all tables and sequences of a group with:

```
SELECT emaj.emaj_snap_group('<group.name>','<storage.directory>');
```

The function returns the number of tables and sequences contained by the group.

This *emaj\_snap\_group()* function generates one file per table and sequence belonging to the supplied tables group. These files are stored in the directory or folder corresponding to the second parameter.

The directory/folder name must be supplied as an absolute pathname and must have been created previously. This directory/folder must have the appropriate permission so that the PostgreSQL cluster can write in it. New files with the same name as existing files will overwrite them.

Created files name has the following pattern:

```
<schema.name>_<table/sequence.name>.snap
```

Each file corresponding to a sequence has only one row, that contains all characteristics of the sequence.

Files corresponding to tables contain one record per row, with the default format of the *COPY* command. These records are sorted in ascending order of the primary key.

All these files are CSV files so that they can be used easily processed by all kind of tools.

At the end of the operation a file named `_INFO` is created in this same directory/folder. It contains a message including the tables group name and the date and time of the snap operation.

It is not necessary that the tables group be in idle state to snap tables.

As this function may generate large or very large files (of course depending on tables sizes), it is user's responsibility to provide a sufficient disk space.

Thanks to this function, a simple test of the E-Maj behaviour could chain:

- `emaj_create_group()`,
- `emaj_start_group()`,
- `emaj_snap_group(<directory_1>)`,
- updates of application tables,
- `emaj_rollback_group()`,
- `emaj_snap_group(<directory_2>)`,
- comparison of both directories content, using a `diff` command for instance.

#### 4.6.7 Snap log tables of a group

It is also possible to record a full or a partial image of all log tables related to a group. This provide a way to archive updates performed by one or several processings. It is possible to dump on files all tables and sequences of a group with:

```
SELECT  
emaj.emaj_snap_log_group('<group.name>', '<start.mark>', '<end.mark>', '<storage.directory>');
```

A `NULL` value or an empty string may be used as start mark. They represents then the first known mark.

The keyword `'EMAJ_LAST_MARK'` can be used as mark name. It then represents the last set mark.

The function returns the number of tables and sequences contained by the group.

This `emaj_snap_log_group()` function generates one file per log table, containing the part of this table that correspond to the updates performed between both supplied marks. Created files name has the following pattern:

`<schema.name>_<table/sequence.name>_log.snap`

The function also generates two files, containing the application sequences state at the time of the respective supplied marks, and named:

`<group.name>_sequences_at_<mark.name>`

These files are stored in the directory or folder corresponding to the last parameter.



The directory/folder name must be supplied as an absolute pathname and must have been created previously. This directory/folder must have the appropriate permission so that the PostgreSQL cluster can write in it. New files with the same name as existing files will overwrite them.

All these files are CSV files so that they can be used easily processed by all kind of tools.

At the end of the operation a file named *\_INFO* is created in this same directory/folder. It contains a message including the tables group name, the marks name that defined the mark range and the date and time of the snap operation.

It is not necessary that the tables group be in idle state to snap log tables.

As this function may generate large or very large files (of course depending on tables sizes), it is user's responsibility to provide a sufficient disk space.

The structure of log tables is directly derived from the structure of their related application table. They contain the same columns with the same type. But they also have some additional technical columns:

- *emaj\_verb* type of the SQL verb that generated the update (INS, UPD, DEL)
- *emaj\_tuple* row version (OLD for DEL and UPD, NEW for INS and UPD)
- *emaj\_id* log row identifier
- *emaj\_changed* log row insertion timestamp
- *emaj\_txid* transaction id that performed the update
- *emaj\_user* connection role that performed the update
- *emaj\_user\_ip* ip address of the client that performed the update (if the client was connected with ip protocol)

## 4.7 PARALLEL ROLLBACK

---

On servers having several processors or processor cores, it may be interesting to reduce rollback elapse time by paralleling the operation on several corridors. For this purpose, E-Maj delivers a specific client to run as a command. It activates E-Maj rollback functions though several parallel connections to the database.

### 4.7.1 Sessions

To run a rollback in parallel, E-Maj spreads tables and sequences to process for one or several tables groups into « sessions ». Each session is then processed in each own corridor.

However, in order to guarantee the integrity of the global operation, the rollback of all sessions is executed inside a single transaction.

To built most balanced sessions as possible, E-Maj takes into account:

- the number of sessions specified by the user in its command,
- statistics about rows to rollback, as reported by the *emaj\_log\_stat\_group()* function,
- foreign key constraints that link several tables between them, 2 updated tables linked by a foreign key constraint being affected into the same session.

### 4.7.2 Prerequisites

The command to run parallel rollbacks is written in php. As a consequence, *php* software and its PostgreSQL interface has to be installed on the server that executes the command (which is not necessarily the same as the one that hosts the PostgreSQL cluster).

Rollbacking each session on behalf of a unique transaction implies the use of two phase commit. As a consequence, the *max\_prepared\_transaction* parameter of the *postgres.conf* file must be adjusted. As the default value of this parameter equals 0, it must be modified by specifying a value at least equal to the maximum number of sessions that will be used.

### 4.7.3 Syntax

The command that performs a parallel rollback has the following syntax:

```
emajParallelRollback.php -g <group(s).name> -m <mark> -s <number.of.sessions>  
[OPTIONS]...
```

General options:

- l specifies that the requested rollback is a “*logged rollback*” (see §4.2.7)
- v displays more information about the execution of the processing
- help only displays a command help
- version only displays the software version

Connection options:

- d database to connect to
- h host to connect to
- p ip-port to connect to
- U connection role to use
- W password associated to the role, if needed

To replace some or all these parameters, the usual *PGDATABASE*, *PGPORT*, *PGHOST* and/or *PGUSER* environment variables can be used.

To specify a list of tables groups in the *-g* parameter, separate the name of each group by a comma.

The supplied connection role must be either a superuser or a role having *emaj\_adm* rights.

For safety reasons, it is not recommended to use the *-W* option to supply a password. It is rather advisable to use the *.pgpass* file (see PostgreSQL documentation).

To let the rollback operation work, the tables group or groups must be in logging state. The supplied mark must also correspond to the same point in time for all groups. In other words, this mark must have been set by the same *emaj\_set\_mark\_group()* function call.

The '*EMAJ\_LAST\_MARK*' keyword can be used as mark name, meaning the last set mark.

In order to test the *emajParallelRollback.php* command, the E-Maj extension supplies a test script, *prep-pr.sql*. It prepares an environment with two tables groups containing some tables and sequences, on which some updates have been performed, with intermediate marks. Once this script has been executed under *psql*, the command displayed at the end of the script can be simply run.

#### 4.7.4 Examples

The command:

```
../php/emajParallelRollback.php -d mydb -g myGroup1 -m Mark1 -s 3
```

logs on database mydb et execute a rollback of group myGroup1 to mark Mark1, using 3 parallel sessions.

The command:

```
../php/emajParallelRollback.php -d mydb -g "myGroup1,myGroup2" -m Mark1 -s 3 -l
```

logs on database mydb and execute a *logged rollback* of both groups myGroup1 and myGroup2 to mark Mark1, using 3 parallel sessions.

## 5 MISCELLANEOUS

---

### 5.1 PARAMETERS

---

The E-Maj extension works with some parameters. Those are stored into the *emaj\_param* internal table.

*emaj\_param* table structure is the following:

- *param\_key* keyword identifying the parameter
- *param\_value\_text* parameter value, if its type is text (otherwise NULL)
- *param\_value\_int* parameter value, if its type is integer (otherwise NULL)
- *param\_value\_boolean* parameter value, if its type is boolean (otherwise NULL)
- *param\_value\_interval* parameter value, if its type is time interval (otherwise NULL)

The E-Maj extension installation procedure inserts a single row into *emaj\_param* table. This row, that should not be modified, describes parameter:

- *version* (text) current E-Maj version.

But the E-Maj administrator may insert other rows into *emaj\_param* table to change the default value of some parameters.

Presented in alphabetic order, the existing key values are:

- *avg\_row\_delete\_log\_duration* (interval) default value = 10 µs ; defines the average duration of a log row deletion ; can be modified to better represent the performance of the server that hosts the database (see § 4.5.3).
- *avg\_row\_rollback\_duration* (interval) default value = 100 µs ; defines the average duration of a row rollback ; can be modified to better represent the performance of the server that hosts the database (see § 4.5.3).
- *fixed\_table\_rollback\_duration* (interval) default value = 5 ms ; defines a fixed rollback cost for any table belonging to a group ; can be modified to better represent the performance of the server that hosts the database (see § 4.5.3).
- *fixed\_table\_with\_rollback\_duration* (interval) default value = 2,5 ms ; defines an additional fixed rollback cost for any table that effectively has updates to rollback ; can be modified to better represent the performance of the server that hosts the database (see § 4.5.3).
- *history\_retention* (interval) default value = 1 mois ; it can be adjusted to change the retention delay of rows in the *emaj\_hist* history table (see § 5.3),

Below is an example of SQL statement that defines a retention delay of rows in the history table, different from the default value (1 month):

```
INSERT INTO emaj.emaj_param (param_key, param_value_interval) VALUES  
( 'history_retention', '15 days'::interval);
```

It is also possible to manage parameter values using any graphic tool such as PgAdmin or phpPgAdmin.

## 5.2 CHECKS

---

When a function is executed to start a tables group, to set a mark or to rollback a tables group, E-Maj performs a certain number of checks in order to verify the integrity of the tables group to process.

These tables group integrity verifications include:

- a check that the PostgreSQL version at tables group creation time is compatible with the current version,
- a check that each application sequence or table of the group always exists,
- a check that each table of the group has its log table, its log and rollback functions and its triggers,
- a check that the log tables structure always reflects the related application tables structure.

## 5.3 TRACES OF OPERATIONS

---

All operations performed by E-Maj, and that impact in any way a tables group, are traced into a table named *emaj\_hist*.

*emaj\_hist* table structure is the following:

- *hist\_id* serial number identifying a row in this history table
- *hist\_datetime* recording date and time of the row
- *hist\_function* function associated to the traced event
- *hist\_event* event kind
- *hist\_object* object name related to the event (group, table or sequence)
- *hist\_wording* additional comments
- *hist\_user* role whose action has generated the event
- *hist\_txid* identifier of the transaction that has generated the event

The *hist\_function* column can take the following values:

- *EMAJ\_INSTALL* E-Maj installation or version update
- *CREATE\_GROUP* tables group creation
- *COMMENT\_GROUP* comment set on a group
- *DROP\_GROUP* tables group suppression
- *FORCE\_DROP\_GROUP* tables group forced suppression
- *START\_GROUP(S)* tables group start
- *STOP\_GROUP(S)* tables group stop
- *LOCK\_GROUP(S)* lock set on tables of a group
- *LOCK\_SESSION(S)* lock set on tables of a session
- *SET\_MARK\_GROUP(S)* mark set on a tables group
- *COMMENT\_MARK\_GROUP* comment set on a mark for a tables group
- *DELETE\_MARK\_GROUP* mark deletion for a tables group
- *RENAME\_MARK\_GROUP* mark rename for a tables group
- *ROLLBACK\_GROUP(S)* rollback updates for a tables group
- *RESET\_GROUP* log tables content reset for a group
- *ROLLBACK\_TABLE* rollback updates for one table
- *ROLLBACK\_SEQUENCE* rollback one sequence
- *SNAP\_GROUP* snap all tables and sequences for a group
- *SNAP\_LOG\_GROUP* snap all log tables for a group

the *hist\_event* column can take the following values:

- *BEGIN*
- *END*
- *MARK DELETED*

*emaj\_hist* content can be viewed by anyone who has the proper access rights on this table (superuser, *emaj\_adm* or *emaj\_viewer* roles).

When a tables group is started (*emaj\_start\_group()* function), the oldest rows are deleted from *emaj\_hist* tables. By default, the retention delay for history rows equals 1 month. But

this value can be modified at any time by inserting the *history\_retention* parameter into *emaj\_param* table with a SQL statement (see § 5.1).

## 5.4 IMPACTS ON CLUSTER AND DATABASE ADMINISTRATION

---

### 5.4.1 Stopping and restarting the cluster

Using E-Maj doesn't bring any particular constraint regarding stopping and restarting a PostgreSQL cluster.

#### 5.4.1.1 General rule

At cluster restart, all E-Maj objects are in the same state as at cluster stop: log triggers of active tables groups remains enabled and log tables contain cancel-able updates already recorded.

If a transaction with table updates were not committed at cluster stop, it would be rolled back during the recovery phase of the cluster start, the application tables updates and the log tables updates being cancelled at the same time.

This rule also applies of course to transactions that execute E-Maj functions, like a tables group start or stop, a rollback, a mark deletion,...

#### 5.4.1.2 Sequences rollback

Due to a PostgreSQL constraint, the rollback of application sequences assigned to a tables group is the only operation that is not protected by transactions. That is the reason for application sequences are processed at the very end end the rollback operations (See §4.2.6)

In case of a cluster stop during an E-Maj rollback execution, it is recommended to rerun this rollback just after the cluster restart, to ensure that application sequences and tables remain properly synchronised.

### 5.4.2 Saving and restoring the database



Using E-Maj may lead to reduce the database saves frequency. But E-Maj cannot be considered as a substitute to regular database saves that remain indispensable to keep a full image of PostgreSQL cluster files on an external support.



### 5.4.2.1 File level saves and restores

When saving or restoring clusters at file level, it is essential to save or restore ALL cluster files. This includes of course all files from the *tspemaj* tablespace.

After a file level restore, tables groups are in the very same state as at the save time, and the database activity can be restarted without any particular E-Maj operation.

### 5.4.2.2 Logical saves and restores of entire database

Regarding stopped tables groups (in *IDLE* state), as log triggers are disabled and the content of related log tables is meaningless, there is nothing to take care to find them in the same state as at save time.

Concerning tables groups in *LOGGING* state at save time, it is important to be sure that log triggers will only be activated after the application tables rebuild. Otherwise, during the tables rebuild, tables updates would also be recorded in log tables!

When using *pg\_dump* command for saves and *psql* or *pg\_restore* commands for restores, and processing full databases (schema + data), these tools recreate triggers, E-Maj log triggers among them, after tables have been rebuilt. So there is no specific precaution to take.

On the other hand, in case of data only save or restore (i.e. without schema, using *-a* or *--data-only* options), the *--disable-triggers* must be supplied:

- with *pg\_dump* (or *pg\_dumpall*) with save in *plain* format (and *psql* is used to restore),
- with *pg\_restore* command with save in *tar* or *custom* format.

### 5.4.2.3 Logical save and restore of partial database

With *pg\_dump* and *pg\_restore* tools, database administrators can perform a subset of database schemas or tables

Restoring a subset of application tables and/or log tables generates a heavy risk of data corruption in case of later E-Maj rollback of concerned tables. Indeed, it is impossible to guarantee in this case that application tables, log tables and internal E-Maj tables that contain essential data for rollback, remain coherent.

If it is necessary to perform partial application tables restores, a drop and recreation of all tables groups concerned by the operation must be performed just after.

The same way, it is strongly recommended to NOT restore a partial *emaj* schema content.

The only case of safe partial restore concerns a full restore of the *emaj* schema content as well as all tables belonging to all groups that are created in the database.

### 5.4.3 Tables reorganisation

No table from *emaj* schema has “cluster” index. So using E-Maj has no operational impact regarding the execution of *CLUSTER* SQL commands at database level.

### 5.4.4 Using E-Maj with replication

E-Maj est totally compatible with the use of the different PostgreSQL replication modes (WAL archiving and PITR, asynchronous and synchronous Streaming Replication). Indeed, all E-Maj objects hosted in the cluster are replicated among all other objects of the cluster.

Using E-Maj with external replication solutions based on triggers like Slony or Londiste, requires some attention... It is probably advisable to avoid to replicate log tables and E-Maj technical tables.

### 5.4.5 PostgreSQL version upgrade

#### 5.4.5.1 PostgreSQL versions 8.2 and 8.3

Tables groups that are created in PostgreSQL version 8.2 or 8.3 can only be managed in their creation version. Indeed, in these PostgreSQL versions, the E-Maj functions behave not totally similarly.

For this reason, when upgrading from 8.2 or 8.3 to a higher version, it is necessary to uninstall and then reinstall E-Maj (see §3.3 and §3.2). As a consequence, it is not possible to keep tables groups in logging state while migrating the PostgreSQL version.

#### 5.4.5.2 PostgreSQL versions 8.4 and later

For all PostgreSQL version greater or equal to 8.4, E-Maj objects and functions are identical.

So it is possible to upgrade the PostgreSQL version without E-Maj reinstallation. Tables groups may remain in logging state at PostgreSQL upgrade.

However, it is recommended to stop tables groups before the PostgreSQL upgrade, tables being normally in a stable state at that time. Moreover, if the PostgreSQL version upgrade is performed using a database dump and restore, the execution of an *emaj\_reset\_group()* function may reduce the volume of data to manipulate, thus reducing the time needed for the operation.

## 5.5 SENSITIVITY TO SYSTEM TIME CHANGE

---

To ensure the integrity of tables managed by E-Maj, it is important that the rollback mechanism be insensitive to potential date or time change of the server that hosts PostgreSQL cluster.

Date and time of each update or each mark is recorded. But nothing but sequence values recorded when marks are set, are used to frame operation in time. So rollbacks and mark deletions are insensitive to potential system date or time change.

However, two minor actions may be influenced by a system date or time change:

- the deletion of oldest events in the *emaj\_hist* table (the retention delay is a time interval),
- finding the name of the mark immediately preceding a given date and time as delivered by the *emaj\_get\_previous\_mark\_group()* function.

## 5.6 USAGE LIMITS

---

The E-Maj extension usage has some limits.

- The minimum required PostgreSQL version is 8.2.
- All tables belonging to a “rollbackable” tables group must have an explicit *PRIMARY KEY*.
- For a table declared in a group, the sum of the schema name length and the table name length can not exceed 52 characters.
- The maximum number of rows that a log table can contain is not unlimited. But this limit is very high (4.611.686.018.427.387.903 !).
- The schema named "*emaj*" is created at E-Maj initialisation. If its name should be changed, the *emaj.sql* scripts, as well as test scripts and the *emajParallelRollback.php* command should be adapted consequently.
- The tablespace named "*tspemaj*" and containing all tables from *emaj* schema must be created before to install the extension. If its name should be changed, *emaj.sql* script should be adapted consequently.
- If a *TRUNCATE* SQL verb is executed on an application table belonging to a group, E-Maj is not able to reset this table in a previous state. Indeed, when a *TRUNCATE* is executed, no trigger is executed at each row deletion. Starting from 8.4 PostgreSQL version, a trigger, created by E-Maj, blocks any *TRUNCATE* statement on any table belonging to a tables group in logging state. For older PostgreSQL versions, this detection is not possible.
- If a DDL operation is executed on an application table belonging to a tables group, E-Maj is not able to reset the table in its previous state.

To understand this last point, it may be interesting to understand the consequences of a DDL operation on the way E-Maj works, depending on the kind of executed operation.

- If a new table were created, it would be able to enter into a group's definition until this group be stopped, dropped and then recreated.
- If a table belonging to a group in logging state were dropped, there would be no way for E-Maj to recover it's structure and its content.
- For a table belonging to a tables group in logging state, adding or deleting a column would generate an error at the next *INSERT/UPDATE/DELETE* SQL verb execution.
- For a table belonging to a tables group in logging state, renaming a column would not generate any error at further log recording. But the checks that E-Maj performs would block any attempt to rollback the related group.
- For a table belonging to a tables group in logging state, changing the type of a column would lead to an inconsistency between the application table and the log table. But, depending on the change of data type applied, updates logging could either work or not. Furthermore, data could be corrupted, for instance in case of increased data length not propagated in log tables. Anyway, due to the checks performed by E-Maj, any attempt to rollback the related group would then fail.
- However, it is possible to create, modify or drop indexes, rights or constraints for a table belonging to a tables group in logging state. But of course, cancelling these changes could not be done by E-Maj.

## 5.7 USER'S RESPONSABILITY

---

### 5.7.1 Defining tables groups content

Defining the content of tables group is essential to guarantee the database integrity. It is the E-Maj administrator's responsibility to ensure that all tables updated by a given processing are really included in a single tables group.

### 5.7.2 Appropriate call of main functions

*emaj\_start\_group()*, *emaj\_set\_mark\_group()*, *emaj\_rollback\_group()* and *emaj\_rollback\_and\_stop\_group()* functions set explicit locks on tables of the group to be sure that no transactions updating these tables are running at the same time. But it is the user's responsibility to execute these operations "at the right time", i.e. at moments that really correspond to stable point in the life of these tables.

### 5.7.3 Management of application triggers

Triggers may have been created on application tables. It is not rare that these triggers performs one or several updates on other tables. In such a case, it is the E-Maj administrator's responsibility to understand the impact of rollback operations on tables concerned by triggers, and if needed to take the appropriate measures.

If the trigger simply adjust the content of the row to insert or update, the logged data will contain the final value of columns. So the rollback would reset the old values without any problem. But may be it would be necessary to deactivate such a trigger during a rollback operation.

If the trigger updates another table, two cases must be considered:

- if the updated table belong to the same tables group, it would be necessary to deactivate the trigger during a rollback operation, so that E-Maj and only E-Maj performs the updates required by the rollback operation,
- if the updated table does not belong to the same tables group, it is essential to analyse the consequences of a rollback operation, in order to avoid a de-synchronisation between both tables. In such a case, only deactivating the trigger may not be sufficient.

#### 5.7.4 Internal E-Maj table or sequence change

With the rights they have been granted, *emaj\_adm* roles and super-users can update any E-Maj internal table.



But in practice, only two tables may be updated by these users: *emaj\_group\_def* and *emaj\_param*. Any other internal table or sequence update may lead to data corruption during rollback operations.

## 6 PHPPGADMIN PLUGIN

---

To make E-Maj use easier, a plugin for the phpPgAdmin 5 administration tool is also available.

### 6.1 GENERAL PRESENTATION

---

For databases into which the E-Maj extension has been installed, and if the user is connected with a role that owns the required rights, E-Maj objects are accessible. It is then possible to:

- see the list of tables groups and perform any possible action, depending on groups state (start, stop, set or remove a mark, rollback, add or modify a comment),
- define or modify groups composition,
- see the list of the marks that have been set for a group, and perform any possible action on them (delete, rename, rollback, add or modify a comment),
- get statistics about log tables content.

### 6.2 USING PHPPGADMIN PLUGIN

---

#### 6.2.1 How to reach E-Maj from phpPgAdmin interface

On figure 1 below, once being connected with a proper role to a database where the E-Maj extension has been installed, a new icon is easily visible in the horizontal database icons tab. Obviously, the *emaj* schema appears in schemas list.

In the browser tree on the left, a new E-Maj object also appears. By opening it, the list of created tables groups becomes directly accessible.

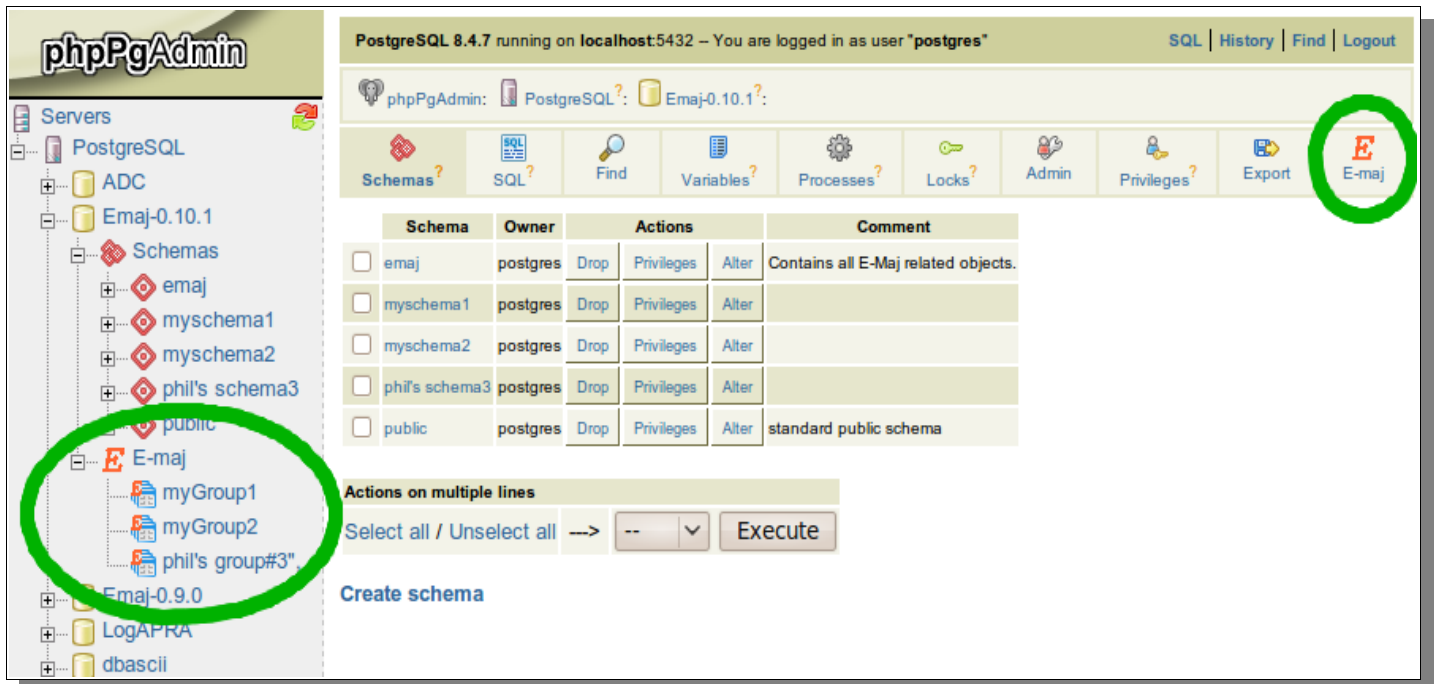


Figure 1 – Connection to a database where E-Maj is installed.

## 6.2.2 Tables groups list

By clicking on one of the E-Maj icons, the user reaches a page that lists all tables groups created in this database (see figure 2 above).

In fact, this page displays two lists: one for groups in LOGGING state and the other for groups in IDLE state.

If some groups are known in the *emaj\_group\_def* table but have not been yet created, the last part of the page allows the user to create tables groups. With a link, the user can also define or modify the composition of tables groups.

All pages displayed by the E-Maj plugin have a header that contains:

- a button to refresh the current page,
- the time at current page display,
- the E-Maj version installed on the database, a click on this field verifying the state of the E-Maj environment,
- the disk space used by the E-Maj tablespace and the part of the whole database it represents,
- the page title.



PostgreSQL 8.4.7 running on localhost:5432 – You are logged in as user "postgres" SQL | History | Find | Logout

phpPgAdmin: PostgreSQL?: Emaj-0.10.1?:

Schemas? SQL? Find Variables? Processes? Locks? Admin Privileges? Export E-maj

16:12:24 E-Maj 0.10.1 [ 472 kB = 5,4% ] - Tables groups list

**Tables groups in "LOGGING" state :**

|                          | Group    | Creation date/time  | # tables | # sequences | Type          | Log size | Actions |      |            |          |               | Comment |
|--------------------------|----------|---------------------|----------|-------------|---------------|----------|---------|------|------------|----------|---------------|---------|
| <input type="checkbox"/> | myGroup1 | 20/12/2011 19:17:35 | 5        | 1           | ROLLBACK-ABLE | 96 kB    | Detail  | Stop | Set a mark | Rollback | Set a comment |         |
| <input type="checkbox"/> | myGroup2 | 20/12/2011 21:36:33 | 4        | 2           | ROLLBACK-ABLE | 64 kB    | Detail  | Stop | Set a mark | Rollback | Set a comment |         |

Actions on multiple lines

Select all / Unselect all → Set a mark ▾ Execute

---

**Tables groups in "IDLE" state :**

|                          | Group            | Creation date/time  | # tables | # sequences | Type       | Log size | Actions |       |       |      |               | Comment |
|--------------------------|------------------|---------------------|----------|-------------|------------|----------|---------|-------|-------|------|---------------|---------|
| <input type="checkbox"/> | phil's group#3", | 20/12/2011 22:10:59 | 2        | 1           | AUDIT-ONLY | 32 kB    | Detail  | Start | Reset | Drop | Set a comment |         |

Actions on multiple lines

Select all / Unselect all → Start ▾ Execute

---

Creation of a new tables group

dummyGrp1 ▾ Create

[Define groups' content](#)

*Figure 2 – List of the tables groups created in the database, with 2 groups in logging state and 1 group in idle state.*

For each tables group, are displayed the following attributes:

- its creation date and time,
- the number of application tables and sequences it contains,
- its type (« ROLLBACK-ABLE » or « AUDIT\_ONLY »),
- the disk space used by its log tables,
- its associated comment, if any.

Several buttons are available so that the user can perform any possible action, depending on the group state.

Under each list, a combo box and a button are dedicated to multi-groups actions.

## 6.2.3 Tables group composition

With a click on the link located at the bottom of tables groups list page, the user reach the function that creates or modifies the tables groups composition.

The upper part of the page lists the existing schemas (except *emaj* schema). By selecting a schema, the list of its tables and sequences appears.

PostgreSQL 8.4.7 running on localhost:5432 – You are logged in as user "postgres"
SQL | History | Find | Logout

phpPgAdmin: PostgreSQL? Emaj-0.10.1?

Schemas?
SQL?
Find
Variables?
Processes?
Locks?
Admin
Privileges?
Export
E-maj

16:12:03 E-Maj 0.10.1 [ 472 kB = 5,4% ] - Tables groups' setup

[Back to the tables groups list](#)

**Schemas list**

| Schema                         | Owner    | Comment                |
|--------------------------------|----------|------------------------|
| <a href="#">myschema1</a>      | postgres |                        |
| <a href="#">myschema2</a>      | postgres |                        |
| <a href="#">phil's schema3</a> | postgres |                        |
| <a href="#">public</a>         | postgres | standard public schema |

---

**Tables et séquences du schéma "myschema1"**

|                          | Type     | Table / Sequence  | Group     | Priority | Actions               | Owner    | Tablespace | Comment |
|--------------------------|----------|-------------------|-----------|----------|-----------------------|----------|------------|---------|
| <input type="checkbox"/> | Table    | myTbl3            | myGroup 1 | 10       | Remove from the group | postgres |            |         |
| <input type="checkbox"/> | Sequence | myTbl3_col31_seq  | myGroup 1 | 1        | Remove from the group | postgres |            |         |
| <input type="checkbox"/> | Table    | mytbl1            | dummyGrp3 |          | Remove from the group | postgres |            |         |
| <input type="checkbox"/> | Table    | mytbl1            | myGroup 1 | 20       | Remove from the group | postgres |            |         |
| <input type="checkbox"/> | Table    | mytbl2            | myGroup 1 |          | Remove from the group | postgres |            |         |
| <input type="checkbox"/> | Table    | mytbl2b           | myGroup 1 |          | Remove from the group | postgres |            |         |
| <input type="checkbox"/> | Sequence | mytbl2b_col20_seq |           |          | Assign to a group     | postgres |            |         |
| <input type="checkbox"/> | Table    | mytbl4            | myGroup 1 | 20       | Remove from the group | postgres |            |         |

**Actions on multiple lines**

Select all / Unselect all → -- v Execute

[Back to the tables groups list](#) | [back to top](#) back to top

Figure 3 – Tables groups composition.

The user can then view or modify the content of the *emaj\_group\_def* table used for the tables groups creation (*emaj\_create\_group()* function).

Are listed for each table or sequence:

- the tables group it belongs to, if any,
- the priority level in the group, if any,
- the identity of its owner,
- the tablespace it belongs to, if any
- the associated comment in the database.

With a button, it is possible to either assign the table or the sequence to a new or already known tables group, or detach it from the tables group it is already assigned.

Note that any change applied in the *emaj\_group\_def* table's content will only take effect when the concerned tables groups will be dropped or created.

#### **6.2.4 Tables group details**

From the tables groups list page, it is possible to get more information about a particular tables group by clicking on its name or on its « Detail » button.

PostgreSQL 8.4.7 running on localhost:5432 – You are logged in as user "postgres" [SQL](#) | [History](#) | [Find](#) | [Logout](#)

phpPgAdmin: PostgreSQL?: Emaj-0.10.1?:

Schemas? SQL? Find Variables? Processes? Locks? Admin Privileges? Export E-maj

16:12:45 E-Maj 0.10.1 [ 472 kB = 5,4% ] - Tables group's detail

[Back to the tables groups list](#) | [Go to bottom](#)

Tables group « myGroup1 » characteristics:

| State   | Creation date/time              | # tables | # sequences | Type          | Log size | Actions |            |         |               |
|---------|---------------------------------|----------|-------------|---------------|----------|---------|------------|---------|---------------|
| Logging | 2011-12-20<br>19:17:35.80117+01 | 5        | 1           | ROLLBACK-ABLE | 96 kB    | Stop    | Set a mark | Content | Set a comment |

Tables group « myGroup1 » marks:

| Mark | Date/Time                        | State  | # row updates | Cumulative updates | Actions        |          |        |        |            |               | Comment |
|------|----------------------------------|--------|---------------|--------------------|----------------|----------|--------|--------|------------|---------------|---------|
| m3   | 2011-12-23<br>18:12:31.27255+01  | Active | 0             | 0                  | Rollback Stats | Rollback | Rename | Delete | First mark | Set a comment |         |
| m2   | 2011-12-23<br>18:01:29.550896+01 | Active | 3             | 3                  | Rollback Stats | Rollback | Rename | Delete | First mark | Set a comment |         |
| m1   | 2011-12-20<br>21:37:55.115042+01 | Active | 1             | 4                  | Rollback Stats | Rollback | Rename | Delete |            | Set a comment |         |

Statistics:

Range start:  Range end:

[Back to the tables groups list](#) | [back to top](#)

Figure 4 – Details of a tables group

A first line repeats information already displayed on the groups list: number of tables and sequences, type, state and disk space used by its log tables.

This line is followed by the group's comment if any has been stored for this group.

A third line presents the possible actions to perform on this group, depending on its state.

Then, the user can see a list of all marks that have been set on the group. For each of them, are displayed:

- its name,
- the date and time it has been set,
- its state,

- the number of recorded log rows between this mark and the next one (or the current situation if this is the last set mark),
- the total number of recorded log rows from when the mark has been set,
- the comment associated to the mark if it exists.

Several buttons are available to perform the actions permitted by its state.

A last line brings a way to get either global or detailed statistics between 2 marks or between a mark and the current situation.

## 6.2.5 Rollback statistics

The figure below shows a page obtained by clicking on the « Rollback Stat » button of a mark.

The screenshot shows the phpPgAdmin interface for PostgreSQL 8.4.7. The top navigation bar includes links for SQL, History, Find, and Logout. Below the navigation bar, there are several utility buttons: Schemas, SQL, Find, Variables, Processes, Locks, Admin, Privileges, Export, and E-maj. The main content area displays the following information:

16:12:15 E-Maj 0.10.1 [ 472 kB = 5,4% ] - Statistics from E-Maj log

[Back to the tables group's detail](#) | [Back to the tables groups list](#)

Rollback tables group « myGroup1 » to mark « m1 » would cancel 4 updates, over 1 table(s), and would take about 00:00:01.

| Schema    | Table  | # row updates |
|-----------|--------|---------------|
| myschema1 | mytbl1 | 4             |

[Back to the tables group's detail](#) | [Back to the tables groups list](#)

*Figure 5 – Before rollback statistics*

The displayed page contains a first line returning the number of log rows that would be concerned by a potential rollback to this mark and an estimate of the rollback duration.

Then, the number of log rows to processed are presented on a per table basis.

## 6.2.6 Statistics about log tables content

These global or detailed statistics are obtained from the detail page of a tables group.

PostgreSQL 8.4.7 running on localhost:5432 – You are logged in as user "postgres" [SQL](#) | [History](#) | [Find](#) | [Logout](#)

phpPgAdmin: PostgreSQL?: Emaj-0.10.1?: myschema1?:

Schemas? SQL? Find Variables? Processes? Locks? Admin Privileges? Export E-maj

16:12:06 E-Maj 0.10.1 [ 472 kB = 5,4% ] - Statistics from E-Maj log

[Back to the tables group's detail](#) | [Back to the tables groups list](#)

Table updates between mark m1 and mark m3 for tables group « myGroup1 »:

| Schema    | Table  | # row updates | Log examination SQL   |
|-----------|--------|---------------|---|
| myschema1 | mytbl1 | 4             | select * from "emaj"."myschema1_mytbl1_log" where emaj_id >= 1 and emaj_id < 9 order by emaj_id |

[Back to the tables group's detail](#) | [Back to the tables groups list](#)

Figure 6 – Global statistics with updates only concerning one application table.

The displayed statistics contain:

- table identity (schema and table names),
- the number of updates recorded into the log tables between both supplied marks,
- a SQL statement that can be copied/pasted to examine the log table's content corresponding to this time frame.

For detailed statistics, there are two additional columns:

- the connection role who performed the updates,
- the SQL verb type (*insert/update/delete*).