

# E-Maj

**A PostgreSQL extension**

**Reference Guide**

**Version: 3.1.0**

# Summary

<b>1 - Introduction</b>	<b>7</b>
1.1 - <a href="#">Document content</a>	7
1.2 - <a href="#">License</a>	7
1.3 - <a href="#">E-Maj's objectives</a>	7
1.4 - <a href="#">Main components</a>	8
<b>2 - How E-Maj works</b>	<b>9</b>
2.1 - <a href="#">Concepts</a>	9
2.1.1 - <a href="#">Tables Group</a>	9
2.1.2 - <a href="#">Mark</a>	9
2.1.3 - <a href="#">Rollback</a>	9
2.2 - <a href="#">Architecture</a>	10
2.2.1 - <a href="#">Logged SQL statements</a>	10
2.2.2 - <a href="#">Created objects</a>	10
2.2.3 - <a href="#">Schemas</a>	11
2.2.4 - <a href="#">Norm for E-Maj objects naming</a>	12
2.2.5 - <a href="#">Tablespaces</a>	12
<b>3 - How to install E-Maj</b>	<b>13</b>
3.1 - <a href="#">Quick start</a>	13
3.1.1 - <a href="#">Software install</a>	13
3.1.2 - <a href="#">Extension install</a>	13
3.1.3 - <a href="#">Extension use</a>	14
3.2 - <a href="#">Installing the E-Maj software</a>	15
3.2.1 - <a href="#">Downloading sources</a>	15
3.2.2 - <a href="#">Standart installation on Linux</a>	15
3.2.3 - <a href="#">Minimum installation on Linux</a>	16
3.2.4 - <a href="#">Installation on Windows</a>	16
3.2.5 - <a href="#">Alternate location of SQL installation scripts</a>	16
3.3 - <a href="#">E-Maj extension setup</a>	17
3.3.1 - <a href="#">Optional preliminary operation</a>	17
3.3.2 - <a href="#">Standart creation of the emaj EXTENSION</a>	17
3.3.3 - <a href="#">Creating the extension by script</a>	18
3.3.4 - <a href="#">Changes in postgresql.conf configuration file</a>	18
3.3.5 - <a href="#">E-Maj parameters</a>	19
3.4 - <a href="#">Update an existing E-Maj version</a>	20
3.4.1 - <a href="#">General approach</a>	20

3.4.2	- <a href="#">Upgrade by deletion and re-installation</a> .....	20
3.4.3	- <a href="#">Upgrade from an E-Maj version between 0.11.0 to 1.3.1</a> .....	22
3.4.4	- <a href="#">E-Maj upgrade from 1.3.1 to a higher version</a> .....	23
3.4.5	- <a href="#">Upgrade an E-Maj version already installed as an extension</a> .....	23
3.5	- <a href="#">Uninstalling an E-Maj extension from a database</a> .....	25
3.6	- <a href="#">PostgreSQL version upgrade</a> .....	26
3.6.1	- <a href="#">Changing PostgreSQL minor versions</a> .....	26
3.6.2	- <a href="#">Changing the major PostgreSQL version and the E-Maj version simultaneously</a> .....	26
3.6.3	- <a href="#">Changing the PostgreSQL major version and keeping the existing E-Maj environment</a> .....	26
3.6.4	- <a href="#">Post migration adaptation script</a> .....	26
<b>4</b>	<b>- <a href="#">How to use E-Maj</a></b> .....	<b>28</b>
4.1	- <a href="#">Set-up the E-Maj access policy</a> .....	28
4.1.1	- <a href="#">E-Maj roles</a> .....	28
4.1.2	- <a href="#">Giving E-Maj rights</a> .....	28
4.1.3	- <a href="#">Giving rights on application tables and objects</a> .....	28
4.1.4	- <a href="#">Synthesis</a> .....	29
4.2	- <a href="#">Main functions</a> .....	30
4.2.1	- <a href="#">Operations chain</a> .....	30
4.2.2	- <a href="#">Define tables groups</a> .....	31
4.2.3	- <a href="#">Create a tables group</a> .....	33
4.2.4	- <a href="#">Start a tables group</a> .....	34
4.2.5	- <a href="#">Set an intermediate mark</a> .....	35
4.2.6	- <a href="#">Rollback a tables group</a> .....	36
4.2.7	- <a href="#">Perform a logged rollback of a tables group</a> .....	38
4.2.8	- <a href="#">Stop a tables group</a> .....	40
4.2.9	- <a href="#">Drop a tables group</a> .....	41
4.3	- <a href="#">Modifying tables groups</a> .....	43
4.3.1	- <a href="#">Modifying a tables group in IDLE state</a> .....	43
4.3.2	- <a href="#">Modifying a tables group in LOGGING state</a> .....	44
4.3.3	- <a href="#">The “Chaining Remove/Add” method</a> .....	47
4.4	- <a href="#">Other groups management functions</a> .....	49
4.4.1	- <a href="#">Reset log tables of a group</a> .....	49
4.4.2	- <a href="#">Comments on groups</a> .....	49
4.4.3	- <a href="#">Protection of a tables group against rollbacks</a> .....	49
4.4.4	- <a href="#">Forced stop of a tables group</a> .....	50
4.4.5	- <a href="#">Forced suppression of a tables group</a> .....	51
4.4.6	- <a href="#">Logged rollback consolidation</a> .....	51
4.4.7	- <a href="#">List of “consolidable rollbacks”</a> .....	52
4.5	- <a href="#">Marks management functions</a> .....	53

4.5.1 - <a href="#">Comments on marks</a> .....	53
4.5.2 - <a href="#">Search a mark</a> .....	53
4.5.3 - <a href="#">Rename a mark</a> .....	54
4.5.4 - <a href="#">Delete a mark</a> .....	54
4.5.5 - <a href="#">Delete oldest marks</a> .....	55
4.5.6 - <a href="#">Protection of a mark against rollbacks</a> .....	56
4.6 - <a href="#">Statistics functions</a> .....	57
4.6.1 - <a href="#">Global statistics about logs</a> .....	57
4.6.2 - <a href="#">Detailed statistics about logs</a> .....	58
4.6.3 - <a href="#">Estimate the rollback duration</a> .....	59
4.7 - <a href="#">Data extraction functions</a> .....	61
4.7.1 - <a href="#">Snap tables of a group</a> .....	61
4.7.2 - <a href="#">Snap log tables of a group</a> .....	62
4.7.3 - <a href="#">SQL script generation to replay logged updates</a> .....	63
4.8 - <a href="#">Other functions</a> .....	66
4.8.1 - <a href="#">Check the consistency of the E-Maj environment</a> .....	66
4.8.2 - <a href="#">Not disabling application triggers at E-Maj rollback time</a> .....	66
4.8.3 - <a href="#">Monitoring rollback operations</a> .....	67
4.8.4 - <a href="#">Updating rollback operations state</a> .....	68
4.8.5 - <a href="#">Deactivating or reactivating event triggers</a> .....	69
4.9 - <a href="#">Multi-groups functions</a> .....	71
4.9.1 - <a href="#">General information</a> .....	71
4.9.2 - <a href="#">Functions list</a> .....	71
4.9.3 - <a href="#">Syntax for groups array</a> .....	71
4.9.4 - <a href="#">Other considerations</a> .....	72
4.10 - <a href="#">Parallel Rollback client</a> .....	73
4.10.1 - <a href="#">Sessions</a> .....	73
4.10.2 - <a href="#">Prerequisites</a> .....	73
4.10.3 - <a href="#">Syntax</a> .....	74
4.10.4 - <a href="#">Examples</a> .....	75
4.11 - <a href="#">Rollback monitoring client</a> .....	76
4.11.1 - <a href="#">Prerequisite</a> .....	76
4.11.2 - <a href="#">Syntax</a> .....	76
- <a href="#">Examples</a> .....	77
<b>5 - <a href="#">Miscellaneous</a></b> .....	<b>78</b>
5.1 - <a href="#">Parameters</a> .....	78
5.2 - <a href="#">Log tables structure</a> .....	79
5.2.1 - <a href="#">Standart structure</a> .....	79
5.2.2 - <a href="#">Adding technical columns</a> .....	79
5.3 - <a href="#">Reliability</a> .....	80
5.3.1 - <a href="#">Internal checks</a> .....	80

5.3.2 - <a href="#">Event triggers</a> .....	81
5.4 - <a href="#">Traces of operations</a> .....	81
5.5 - <a href="#">Impacts on instance and database administration</a> .....	84
5.5.1 - <a href="#">Stopping and restarting the instance</a> .....	84
5.5.2 - <a href="#">Saving and restoring the database</a> .....	84
5.5.3 - <a href="#">Data load</a> .....	86
5.5.4 - <a href="#">Tables reorganisation</a> .....	86
5.5.5 - <a href="#">Using E-Maj with replication</a> .....	87
5.6 - <a href="#">Sensitivity to system time change</a> .....	90
5.7 - <a href="#">Performance</a> .....	91
5.7.1 - <a href="#">Updates recording overhead</a> .....	91
5.7.2 - <a href="#">E-Maj rollback duration</a> .....	91
5.7.3 - <a href="#">Optimizing E-Maj operations</a> .....	91
5.8 - <a href="#">Usage limits</a> .....	93
5.9 - <a href="#">User's responsibility</a> .....	94
5.9.1 - <a href="#">Defining tables groups content</a> .....	94
5.9.2 - <a href="#">Appropriate call of main functions</a> .....	94
5.9.3 - <a href="#">Management of application triggers</a> .....	94
5.9.4 - <a href="#">Internal E-Maj table or sequence change</a> .....	95
<b>6 - <a href="#">Emaj_web</a>.....</b>	<b>96</b>
6.1 - <a href="#">Overview</a> .....	96
6.2 - <a href="#">Emaj_web client installation</a> .....	96
6.2.1 - <a href="#">Prerequisite</a> .....	96
6.2.2 - <a href="#">Download</a> .....	96
6.2.3 - <a href="#">Application configuration</a> .....	96
6.3 - <a href="#">Using Emaj_web</a> .....	97
6.3.1 - <a href="#">Access to Emaj_web and databases</a> .....	97
6.3.2 - <a href="#">Tables groups list</a> .....	98
6.3.3 - <a href="#">Some details about the user interface</a> .....	100
6.3.4 - <a href="#">Tables group details</a> .....	101
6.3.5 - <a href="#">Statistics</a> .....	102
6.3.6 - <a href="#">Tables group content</a> .....	104
6.3.7 - <a href="#">Tables groups configuration</a> .....	105
6.3.8 - <a href="#">Monitoring rollback operations</a> .....	107
6.3.9 - <a href="#">Schemas content</a> .....	108
6.3.10 - <a href="#">E-Maj environment state</a> .....	108
<b>7 - <a href="#">Contribute to the E-Maj development</a>.....</b>	<b>110</b>
7.1 - <a href="#">Build the E-Maj environment</a> .....	110
7.1.1 - <a href="#">Clone the E-Maj repository</a> .....	110
7.1.2 - <a href="#">Description of the E-Maj tree</a> .....	110

7.1.3 - <a href="#">Setting tools parameters</a> .....	110
7.2 - <a href="#">Coding</a> .....	111
7.2.1 - <a href="#">Versionning</a> .....	111
7.2.2 - <a href="#">Coding rules</a> .....	111
7.2.3 - <a href="#">Version upgrade script</a> .....	112
7.3 - <a href="#">Testing</a> .....	113
7.3.1 - <a href="#">Create PostgreSQL instances</a> .....	113
7.3.2 - <a href="#">Install software dependancies</a> .....	113
7.3.3 - <a href="#">Execute non regression tests</a> .....	113
7.3.4 - <a href="#">Test coverage</a> .....	116
7.3.5 - <a href="#">Evaluate the performances</a> .....	116
7.4 - <a href="#">Documenting</a> .....	118
7.5 - <a href="#">Submitting a patch</a> .....	118
<b>8 - <a href="#">Appendix</a>.....</b>	<b>119</b>
8.1 - <a href="#">E-Maj functions list</a> .....	119
8.2 - <a href="#">E-Maj distribution content</a> .....	121
8.3 - <a href="#">PostgreSQL and E-Maj versions compatibility matrix</a> .....	124

# 1 INTRODUCTION

---

## 1.1 DOCUMENT CONTENT

---

This document is a reference 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 E-Maj installation, update and uninstall procedures.

Chapter 4 details how to use E-Maj. It contains a description of each function.

Chapter 5 gives some additional information needed for a good understanding of how the extension works.

Then, chapter 6 presents Emaj\_web, a web graphic interfaces that complement the E-Maj extension.

## 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 ».

It meets two main goals:

- E-Maj can be used to **trace updates** performed by application programs on the table's content. Viewing these recorded updates offers an answer to the need for “updates-auditing”,
- By using these recorded updates, E-Maj is able 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.

In other words, E-Maj is a PostgreSQL extension which enables fine-grained write logging and time travel on subsets of the database.

It provides a good solution to :

- define save points at precise time on a set of tables,
- restore, if needed, this table set into a stable state, without stopping the instance,
- manage several save points, each of them being usable at any time as a 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 database subsets into predefined stable states, so that tests can be replayed as many times as needed.


## 1.4 MAIN COMPONENTS

---

E-Maj actually groups several components:

- a PostgreSQL *extension* object created into each database, named *emaj* and holding some tables, functions, sequences, ...
- a set of external clients working in command line interface,
- a web application,, *Emaj\_web*.

The external clients and the GUI call the functions of the *emaj* extension.



All these components are described in this documentation.



## 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 of a database that are updated by one or more sets of programs. Each tables group is defined by a name which must be unique inside its database. By extent, a tables group can also contain partitions of partitionned tables and sequences. Tables (including partitions) 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 latter case, it is not possible to rollback the group. But using this type of group allows to record tables updates for auditing purposes, even with tables that have no explicitly created primary key or with tables of type *UNLOGGED* or *WITH OIDS*.

#### 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 of resetting all tables and sequences of a group in the state they had when a mark was set.

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* », update cancellations are recorded in log tables, so that they can be later cancelled: the rollback operation can be ... rolled back.

Note that this concept of E-Maj rollback is different from the usual concept of “transactions rollback” managed by PostgreSQL.

## 2.2 ARCHITECTURE

---

In order to be able to perform a rollback operation without having previously kept a physical image of the PostgreSQL instance'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 Logged SQL statements

The recorded update operations concerns the following SQL verbs:

- rows insertions:
  - *INSERT*, either elementary (*INSERT ... VALUES*) or set oriented (*INSERT ... SELECT*)
  - *COPY ... FROM*
- rows updates:
  - *UPDATE*
- rows deletions:
  - *DELETE*
- tables truncations
  - *TRUNCATE*

For statements that process several rows, each creation, update or deletion is individually recorded. For instance, if a “*DELETE FROM <table>*” is performed against a table having 1 million rows, 1 million row deletion events are recorded.

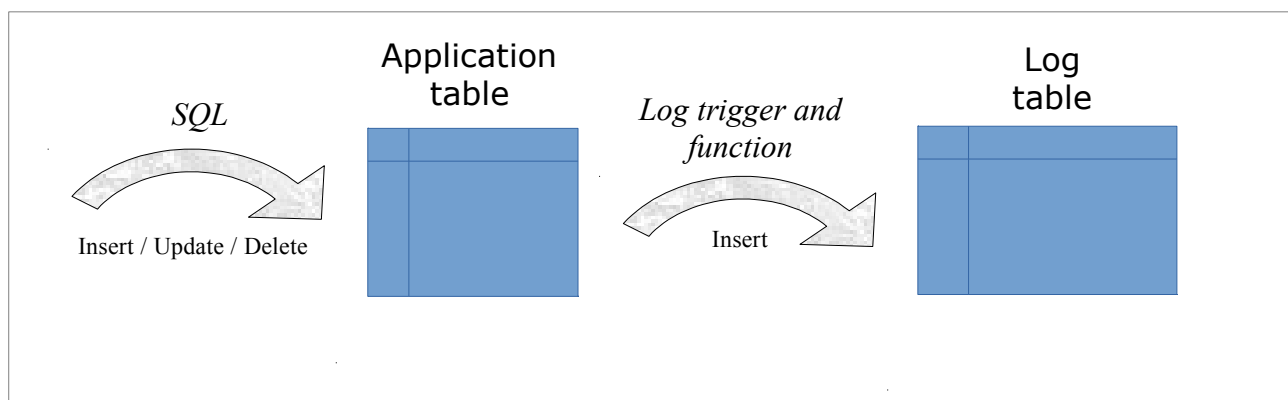
The case of *TRUNCATE* SQL verbs is specific. As no “*FOR EACH ROW*” trigger can be fired for this verb, the consequences of a *TRUNCATE* cannot be cancelled by E-Maj. Therefore, its execution is forbidden for “*rollbackable*” tables groups in “*logging*” state. In contrast, *TRUNCATE* is always permitted for “*audit\_only*” tables groups. In such a case, only its execution is recorded.

### 2.2.2 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*, *COPY*), change (*UPDATE*) or suppression (*DELETE*), record into the log table all data needed to potentially cancel later this elementary action,
- another trigger, that either blocks any execution of a *TRUNCATE* SQL verb for “*rollbackable*” tables groups or records the execution of a *TRUNCATE* SQL verb for “*audit\_only*” tables groups,

- a sequence used to quickly count the number of updates recorded in log tables between 2 marks.



A log table has the same structure as its corresponding application table. However, it contains some additional technical columns, described at §5.2.1.

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

- 17 tables,
- 8 composite and 3 enum types,
- 1 view,
- 2 triggers,
- more than 120 functions, about 50 of them being directly callable by users,
- 1 sequence named *emaj\_global\_seq* used to assign to every update recorded in any log table of the database a unique identifier with an increasing value over time,
- 1 specific schema, named *emaj*, that contains all these relational objects,
- 2 roles acting as groups (NOLOGIN): *emaj\_adm* to manage E-Maj components, and *emaj\_viewer* to only look at E-Maj components
- 3 event triggers.

Technical tables, whose structure is interesting to know, are described in the coming chapters (*emaj\_group\_def* is described in §4.2.2, *emaj\_param* is described in §5.1 and *emaj\_hist* is described in §5.4).

### 2.2.3 Schemas

Almost all technical objects created at E-Maj installation are located into the schema named *emaj*. The only exception is the event trigger « *emaj\_protection\_trg* » that belongs to the *public* schema.

All objects linked to application tables are stored into schemas named :

*emaj\_<schema>*

where <schema> is the schema name of the application tables.

The creation and the suppression of log schemas are only managed by E-Maj functions. They should NOT contain any other objects than those created by the extension.

## 2.2.4 Norm for E-Maj objects naming

For an application table, the log objects name is prefixed with the table name. More precisely, for an application table:

- the name of the log table is:  
*<table.name>\_log*
- the name of the log function is:  
*<table.name>\_log\_fnct*
- the name of the sequence associated to the log table is:  
*<table.name>\_log\_seq*

For application tables whose name is very long (over 50 characters), the prefix used to build the log objects name is generated so it respects the PostgreSQL naming rules and avoids name conflict.

A log table name may contain a suffix like “\_1”, “\_2”, etc. In such a case, it deals with an old log table that has been renamed.

Other E-Maj function names are also normalised:

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

Triggers created on application tables have the same name:

- *emaj\_log\_trg* for the log triggers,
- *emaj\_trunc\_trg* for the triggers that manage TRUNCATE verbs.

The name of event triggers starts with “*emaj\_*” and ends with “*\_trg*”.

## 2.2.5 Tablespaces

When the extension is installed, the E-Maj technical tables are stored into the default tablespace set at instance or database level or explicitly set for the current session.

The same rule applies for log tables and index. But using tables group parameters (Refer to §4.2.2), it is also possible to store log tables and/or their index into specific tablespaces.

## 3 HOW TO INSTALL E-MAJ

---

In this chapter, we will describe how to download and install or upgrade the E-Maj extension. Uninstallation is also discussed in this chapter.

### 3.1 QUICK START

---

The E-Maj installation is described in detail in the coming chapter. But the few following commands allow to quickly install and use E-Maj under Linux.

#### 3.1.1 Software install

To install E-Maj, log on your *postgres* (or another) account, download E-Maj from PGXN (<https://pgxn.org/dist/e-maj/>) and type:

```
unzip e-maj-<version>.zip  
cd e-maj-<version>/  
sudo cp emaj.control $(pg_config --sharedir)/extension/.  
sudo cp sql/emaj--* $(pg_config --sharedir)/extension/.
```

The chapter §3.2 presents other ways to download and install the software.

#### 3.1.2 Extension install

To install the *emaj* extension into a database, log on the target database, using a *superuser* role and execute:

```
create extension emaj cascade;  
grant emaj_adm to <role>;
```

For PostgreSQL versions prior 9.6, see §3.3.2.

With the latest statement, you give E-Maj administration grants to a particular role. Then, this role can be used to execute all E-Maj operations, avoiding the use of *superuser* role.

### 3.1.3 Extension use

You can now log on the database with the role having the E-Maj administration rights.

Then the *emaj\_group\_def* table that defines groups must be populated with one row per table or sequence to link to the group. A table can be added with a statement like:

```
INSERT INTO emaj.emaj_group_def  
      (grpdef_group, grpdef_schema, grpdef_tbldseq)  
VALUES ('my_group', 'my_schema', 'my_table');
```

or select all tables of a schema with:

```
INSERT INTO emaj.emaj_group_def  
      (grpdef_group, grpdef_schema, grpdef_tbldseq)  
SELECT 'my_group', 'my_schema', table_name  
FROM information_schema.tables  
WHERE table_schema = 'my_schema' AND  
      table_type = 'BASE TABLE';
```

knowing that tables inserted into a group must have a primary key.

Then the typical sequence:

```
SELECT emaj.emaj_create_group('my_group');  
SELECT emaj.emaj_start_group('my_group', 'Mark-1');  
[INSERT/UPDATE/DELETE on tables]  
SELECT emaj.emaj_set_mark_group('my_group','Mark-2');  
[INSERT/UPDATE/DELETE on tables]  
SELECT emaj.emaj_set_mark_group('my_group','Mark-3');  
[INSERT/UPDATE/DELETE on tables]  
SELECT emaj.emaj_rollback_group('my_group','Mark-2');  
SELECT emaj.emaj_stop_group('my_group');  
SELECT emaj.emaj_drop_group('my_group');
```

would create and start the tables group, log updates and set several intermediate marks, go back to one of them, stop the recording and finally drop the group.

Additionally, the Emaj\_web can also be installed and used (see §6).

## 3.2 INSTALLING THE E-MAJ SOFTWARE

---

### 3.2.1 Downloading sources

E-Maj is available for download on the Internet site PGXN, the PostgreSQL Extension Network (<https://pgxn.org/dist/e-maj/>).

E-Maj and its add-ons are also available on the github.org Internet site:

- source components (<https://github.com/dalibo/emaj>)
- documentation ([https://github.com/beaud76/emaj\\_doc](https://github.com/beaud76/emaj_doc))
- Emaj\_web GUI ([https://github.com/dalibo/emaj\\_web](https://github.com/dalibo/emaj_web))

For the records, an old repository ([https://github.com/beaud76/emaj\\_ppa\\_plugin](https://github.com/beaud76/emaj_ppa_plugin)) contains a plugin for *phpPgAdmin*. It is not maintained since E-Maj 3.0.

### 3.2.2 Standart installation on Linux

Download the latest E-Maj version by any convenient way. If the *pgxn* client is installed, just execute the command:

```
pgxn download E-Maj
```

Then decompress the downloaded archive file with the commands::

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

Identify the precise location of the *SHAREDIR* directory. Depending on the PostgreSQL installation, the *pg\_config --sharedir* shell command may directly report this directory name. Otherwise, look at typical locations like:

- */usr/share/postgresql/<pg\_version>* for Debian or Ubuntu
- */usr/pgsql-<pg\_version>/share* for RedHat or CentOS

Copy some files to the extension directory of the postgresql version you want to use. As a super-user or pre-pended with *sudo*, type:

```
cp emaj.control <SHAREDIR_directory>/extension/.  
cp sql/emaj--* <SHAREDIR_directory>/extension/.
```

The latest E-Maj version is now installed and referenced by PostgreSQL. The e-maj-<version> directory contains the file tree described at §8.2.

### 3.2.3 Minimum installation on Linux

On some environments (like DBaaS clouds for instance), it is not allowed to add extensions into the *SHAREDIR* directory. For these cases, a minimum installation is possible.

Download the latest E-Maj version by any convenient way and decompress it.

If the *pgxn* client is installed, just execute the commands:

```
pgxn download E-Maj
unzip e-maj-<version>.zip
```

The e-maj-<version> directory contains the file tree described at §8.2.

The extension creation will be a little bit different (§3.3).

### 3.2.4 Installation on Windows

To install E-Maj on Windows:

- Download the extension from the *pgxn.org* site,
- Extract the file tree from the downloaded zip file,
- Copy the files *emaj.control* and *sql/emaj--\** into the *share\extension* folder of the PostgreSQL installation folder (typically *c:\Program\_Files\PostgreSQL\<postgres\_version>*)

### 3.2.5 Alternate location of SQL installation scripts

The *emaj.control* file located in the *SHAREDIR/extension* directory of the PostgreSQL version, may contain a directive that defines the directory where SQL installation scripts are located.

So it is possible to only put the *emaj.control* file into this *SHAREDIR/extension* directory, by creating a pointer towards the script directory.

To setup this, just:

- copy the *emaj.contol* file from the root directory of the decompressed structure into the *SHAREDIR/extension* directory,



➤ adjust the *directory* parameter of the *emaj.control* file to reflect the actual location of the E-Maj SQL scripts.

### 3.3 E-MAJ EXTENSION SETUP

---

If an extension already exists in the database, but in an old E-Maj version, please go on with §3.4.

The standart way to install E-Maj consists in creating an *EXTENSION* object (in the PostgreSQL terminology). To achieve this task, the user must be logged to the database as *superuser*.

In environments for which this is not possible (cases of minimum installation – see §3.2.3), a *psql* script can be executed.

#### 3.3.1 Optional preliminary operation

The technical tables of the E-Maj extension are created into the default tablespace. If the E-Maj administrator wants to store them into a dedicated tablespace, he can create it if needed and define it as the default tablespace for the session:

```
SET default_tablespace = <tablespace.name>;
```

#### 3.3.2 Standart creation of the emaj EXTENSION

##### 3.3.2.1 PostgreSQL version 9.6 and above

The E-Maj extension can now be created into the database, by executing the SQL command:

```
CREATE EXTENSION emaj CASCADE;
```

To start with, the script verifies that the PostgreSQL version is at least 9.5, and that the current user has the *superuser* attribute.

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



*The emaj* schema must only contain E-Maj related objects.

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

Finally, the installation script looks at the instance configuration and may display a warning message regarding the *-max-prepared-statements* parameter (see §4.10.2).

### 3.3.2.2 PostgreSQL version 9.5

With PostgreSQL version prior 9.6, the *CASCADE* clause doesn't exist. In that case, the required extensions must be explicitly created, if needed, before *emaj*.

```
CREATE EXTENSION IF NOT EXISTS dblink;  
CREATE EXTENSION IF NOT EXISTS btree_gist;  
CREATE EXTENSION emaj;
```

### 3.3.3 Creating the extension by script

When creating the *emaj EXTENSION* is not possible, a *psql* script can be used instead.

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

where *<emaj\_directory>* is the directory generated by the E-Maj installation (see §3.2.3) and *<version>* the current E-Maj version.



It is not mandatory to execute the installation script as superuser. But if it is not the case, the role used for this installation will need the rights to create triggers on the application tables of the future tables groups.

### 3.3.4 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. This parameter is used by PostgreSQL to compute the size of the “*shared lock table*” that tracks locks for the whole instance. Its default value equals 64. It can be increased if an E-Maj operation fails with a message indicating that all entries of the “*shared lock table*” have been used.

Furthermore, if the parallel rollback tool may be used (see § 4.10), it will be probably necessary to adjust the ***max\_prepared\_transaction*** parameter.

### 3.3.5 E-Maj parameters

Several parameters have an influence on the E-Maj behaviour. They are presented in details in chapter §5.1.

The parameters setting step is optional. E-Maj works well with the default parameter values.

However, if the E-Maj administrator wishes to take benefit from the rollback operations monitoring capabilities, it is necessary to insert a row into the *emaj\_param* table to setup the value of the “*dblink\_user\_password*” parameter (see §4.8.4.1)

Test and demonstration

It is possible to check whether the E-Maj installation works fine, and discover its main features by executing a demonstration script. Under psql, just execute the *emaj\_demo.sql* script that is supplied with the extension.

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

If no error is encountered, the script displays this final message:

```
### This ends the E-Maj demo. Thank You for using E-Maj and have fun!
```

Examining the messages generated by the script execution, allows to discover most E-Maj features. Once the script execution is completed, the demonstration environment is left as is, so that it remains possible to examine it or to play with it. To suppress it, execute the cleaning function that the script has created.

```
SELECT emaj.emaj_demo_cleanup();
```

This drops the *emaj\_demo\_app\_schema* schema and both *emaj demo group 1* and *emaj demo group 2* tables groups.

## 3.4 UPDATE AN EXISTING E-MAJ VERSION

---

### 3.4.1 General approach

The first step consists in installing the new version of the E-Maj software, using the procedure described in §3.2. Keep the old E-Maj version directory at least until the end of the upgrade. Some files may be needed.

It is also necessary to check whether some preliminary operations as described in §3.3.1 must be executed.

Then, the process to upgrade an E-Maj extension in a database depends on the already installed E-Maj version.

For E-Maj versions prior 0.11.0, there is no particular update procedure. A simple E-Maj deletion and then re-installation has to be done, as described in §3.4.2. This approach can also be used for any E-Maj version, even though it has a drawback: all log contents are deleted, resulting in no further way to rollback or look at the recorded updates.

For installed E-Maj version 0.11.0 and later, it is possible to perform an upgrade without E-Maj deletion.

The upgrade procedure for an existing E-Maj version between 0.11.0 and 1.3.1 is described in §3.4.3.

The upgrade from version 1.3.1 to a higher version is described in §3.4.4.

The upgrade from a version greater or equal 2.0.0 is described in §3.4.5.



Starting from version 2.2.0, E-Maj no longer supports PostgreSQL versions prior 9.2. Starting from version 3.0.0, E-Maj no longer supports PostgreSQL versions prior 9.5. If an older PostgreSQL version is used, it must be updated before migrating E-Maj to a higher version.

### 3.4.2 Upgrade by deletion and re-installation

For this upgrade path, it is not necessary to use the full un-installation procedure described in §3.5. In particular, the tablespace and both roles can remain as is. However, it may be judicious to save some useful pieces of information. Here is a suggested procedure.

#### 3.4.2.1 Stop tables groups

If some tables groups are in *LOGGING* state, they must be stopped, using the *emaj\_stop\_group()* function (see §4.2.8) (or the *emaj\_force\_stop\_group()* function if *emaj\_stop\_group()* (see §4.4.4) returns an error).

### 3.4.2.2 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 instance with a `!copy` command, or by duplicating the table outside the *emaj* schema with a SQL statement like:

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

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

```
CREATE TABLE public.sav_param AS  
SELECT * FROM emaj.emaj_param WHERE param_key <> 'emaj_version';
```

If the installed E-Maj version is 3.1.0 or higher, and if the E-Maj administrator has registered application triggers as “not to be automatically disabled at E-Maj rollback time”, it is advisable to save this triggers list, for instance with:

```
CREATE TABLE public.sav_ignored_app_trigger AS  
SELECT * FROM emaj.emaj_ignored_app_trigger;
```

### 3.4.2.3 E-Maj deletion and re-installation

Once connected as super-user, just chain the execution of the *uninstall.sql* script, of the current version and then the extension creation.

```
\i <old_emaj_directory>/sql/uninstall.sql  
CREATE EXTENSION emaj;
```

NB: starting from E-Maj 2.0.0, the uninstall script is named *emaj\_uninstall.sql*.

### 3.4.2.4 Restore user data

Data previously saved can now be restored into E-Maj both technical tables, for instance with *INSERT ... SELECT* statements.

```

INSERT INTO emaj.emaj_group_def
  SELECT grpdef_group, grpdef_schema, grpdef_tblseq,
         grpdef_priority, grpdef_log_dat_tsp, grpdef_log_idx_tsp
  FROM public.sav_group_def;

INSERT INTO emaj.emaj_param
  SELECT * FROM public.sav_param;

INSERT INTO emaj.emaj_ignored_app_trigger
  SELECT * FROM public.sav_ignored_app_trigger;

```

Once data are copied, temporary tables or files can be deleted.

### 3.4.3 Upgrade from an E-Maj version between 0.11.0 to 1.3.1

For installed version between 0.11.0 and 1.3.1, psql upgrade scripts are supplied. They allow to upgrade from one version to the next one.

Each step can be performed without impact on existing tables groups. They may even remain in *LOGGING* state during the upgrade operations. This means in particular that:

- updates on application tables can continue to be recorded during and after this version change,
- a « *rollback* » on a mark set before the version change can also be performed after the migration.

Source version	Target version	psql script	Duration	Concurrent updates (1)
0.11.0	0.11.1	emaj-0.11.0-to-0.11.1.sql	Very quick	Yes
0.11.1	1.0.0	emaj-0.11.1-to-1.0.0.sql	Very quick	Yes
1.0.0	1.0.1	emaj-1.0.0-to-1.0.1.sql	Very quick	Yes
1.0.1	1.0.2	emaj-1.0.1-to-1.0.2.sql	Very quick	Yes
1.0.2	1.1.0	emaj-1.0.2-to-1.1.0.sql	Variable	No (2)
1.1.0	1.2.0	emaj-1.1.0-to-1.2.0.sql	Very quick	Yes
1.2.0	1.3.0	emaj-1.2.0-to-1.3.0 .sql	Quick	Yes (3)
1.3.0	1.3.1	emaj-1.3.0-to-1.3.1.sql	Very quick	Yes

(1) The last column indicates whether the E-Maj upgrade can be executed while application tables handled by E-Maj are accessed in update mode. Note that any other E-Maj operation executed during the upgrade operation would wait until the end of the upgrade.

(2) When upgrading into 1.1.0, log tables structure changes. As a consequence:

- eventhough tables groups may remain in *LOGGING* state, the upgrade can only be executed during a time period when application tables are not updated by any application processing,
- the operation duration will mostly depends on the volume of data stored into the log tables.

Note also that E-Maj statistics collected during previous rollback operations are not kept (due to large differences in the way rollbacks are performed, the old statistics are not pertinent any more).

(3) It is advisable to perform the upgrade into 1.3.0 in a period of low database activity. This is due to *Access Exclusive* locks that are set on application tables while the E-Maj triggers are renamed.

At the end of each upgrade step, the script displays the following message :

```
>>> E-Maj successfully migrated to <new_version>
```

### 3.4.4 E-Maj upgrade from 1.3.1 to a higher version

The upgrade from the 1.3.1 version is specific as it must handle the installation mode change, moving from a *psql* script to an *extension*.

Concretely, the operation is performed with a single SQL statement:

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

The PostgreSQL extension manager determines the scripts to execute depending on the E-Maj version identifier found in the *emaj.control* file.

But this upgrade is not able to process cases when at least one tables group has been created with a PostgreSQL version prior 8.4. In such a case, these old tables groups must be dropped before the upgrade and recreated after.

### 3.4.5 Upgrade an E-Maj version already installed as an extension

An existing version already installed as an extension can be upgraded using the SQL statement:

```
ALTER EXTENSION emaj UPDATE;
```

The PostgreSQL extension manager determines the scripts to execute depending on the current installed E-Maj version and the version found in the *emaj.control* file.

The operation is very quick et does not alter tables groups. They may remain in *LOGGING* state during the upgrade. As for previous upgrades, this means that:

- updates on application tables can continue to be recorded during and after this version change,
- a « *rollback* » on a mark set before the version change can also be performed after the migration.

Version specific details:

- The procedure that upgrades a version 2.0.1 into 2.1.0, may modify the *emaj\_group\_def* table in order to reflect the fact that the *tspemaj* tablespace is not automatically considered as a default tablespace anymore. If *tspemaj* was effectively used as default tablespace for created tables groups, the related *grpdef\_log\_dat\_tsp* and *grpdef\_log\_idx\_tsp* columns content of the *emaj\_group\_def* table is automatically adjusted so that a future drop and recreate operation would store the log tables and indexes in the same tablespace. The administrator may review these changes to be sure this corresponds to his expectations. (Cf §4.2.2)
- The procedure that upgrades a version 2.2.2 into 2.2.3 checks the recorded log sequences values. In some cases, it may ask for a preliminary reset of some tables groups.
- The procedure that upgrades a version 2.3.1 into 3.0.0 changes the structure of log tables: both *emaj\_client\_ip* and *emaj\_client\_port* columns are not created anymore. Existing log tables are not modified. Only the new log tables are impacted. But the administrator can add these columns, by using the '*alter\_log\_tables*' parameter (Cf §5.2.2).
- The procedure that upgrades a version 3.0.0 into 3.1.0 renames existing log objects. This leads to locking the application tables, which may generate conflicts with the parallel use of these tables. This procedure also issues a warning message indicating that the changes in E-Maj rollback functions regarding the application triggers processing may require changes in user's procedures.



## 3.5 UNINSTALLING AN E-MAJ EXTENSION FROM A DATABASE

---

To uninstall E-Maj from a database, the user must log on this database with `psql`, as a superuser.

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.

Although E-Maj is usually installed with a `CREATE EXTENSION` statement, it cannot be uninstalled with a simple `DROP EXTENSION` statement. An event trigger blocks such a statement.

To uninstall E-Maj, just execute the `emaj_uninstall.sql` delivered script.

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

This script performs the following steps:

- it executes the cleaning functions created by demo or test scripts, if they exist,
- it stops the tables groups in `LOGGING` state, if any,
- it drops the tables groups, removing in particular the triggers on application tables,
- it drops the extension and the main `emaj` schema,
- it drops roles `emaj_adm` and `emaj_viewer` if they are not linked to any objects in the current database or in other databases of the instance.

The uninstallation script execution displays:

```
$ psql ... -f sql/emaj_uninstall.sql  
>>> Starting E-Maj uninstallation procedure...  
SET  
psql:sql/emaj_uninstall.sql:203: WARNING:  emaj_uninstall:  
emaj_adm and emaj_viewer roles have been dropped.  
DO  
SET  
>>> E-maj successfully uninstalled from this database
```

## 3.6 POSTGRESQL VERSION UPGRADE

---

### 3.6.1 Changing PostgreSQL minor versions

As changing the minor PostgreSQL version only consists in replacing the binary files of the software, there is no particular constraint regarding E-Maj.

### 3.6.2 Changing the major PostgreSQL version and the E-Maj version simultaneously

A PostgreSQL major version change may be the opportunity to also change the E-Maj version. But in this case, the E-Maj environment has to be recreated from scratch and old objects (tables groups, logs, marks,...) cannot be reused.

### 3.6.3 Changing the PostgreSQL major version and keeping the existing E-Maj environment

Nevertheless, it is possible to keep the existing E-Maj components (tables groups, logs, marks,...) while changing the PostgreSQL major version. And the tables groups may even stay in logging state during the operation.

But 2 conditions must be met:

- the old and new instances must share the same E-Maj version,
- a post migration script must be executed, before any E-Maj use.

Of course, it is possible to upgrade the E-Maj version before or after the PostgreSQL version change.

If the PostgreSQL version upgrade is performed using a database dump and restore, and if the tables groups may be stopped, a log tables purge, using the *emaj\_reset\_group()* function, may reduce the volume of data to manipulate, thus reducing the time needed for the operation.

### 3.6.4 Post migration adaptation script

It may happen that a PostgreSQL major version change has an impact on the E-Maj extension content. Thus, a script is supplied to handle such cases.

After each PostgreSQL major version upgrade, a *psql* script must be executed as superuser:

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

For E-Maj versions 2.0.0 and later, the script only creates the event triggers that may be missing:

- those that appear in version 9.3 and protect against the drop of the extension itself and the drop of E-Maj objects (log tables, functions,...),
- those that appear in version 9.5 and protect against log table structure changes.

The script may be executed several times on the same version, only the first execution modifying the environment.

## 4 HOW TO USE E-MAJ

---

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

---

A bad usage of E-Maj can break the database 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 reasons, it is preferable to take advantage of both roles created by the installation script:

- *emaj\_adm* is used as the 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>;
```

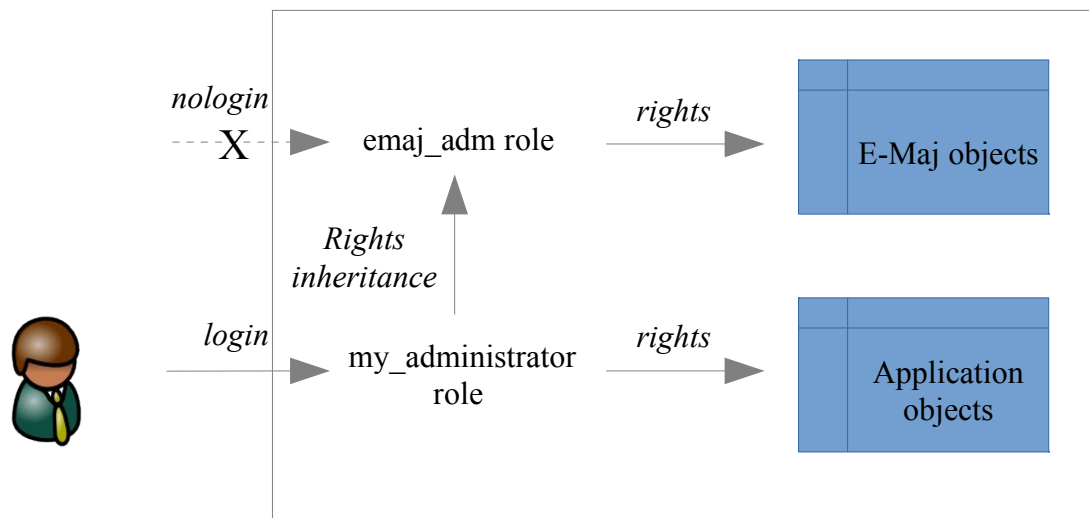
Of course, *emaj\_adm* or *emaj\_viewer* rights can be given to several roles.

#### 4.1.3 Giving rights on application tables and objects

To let an E-Maj administrator also access application tables or 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.

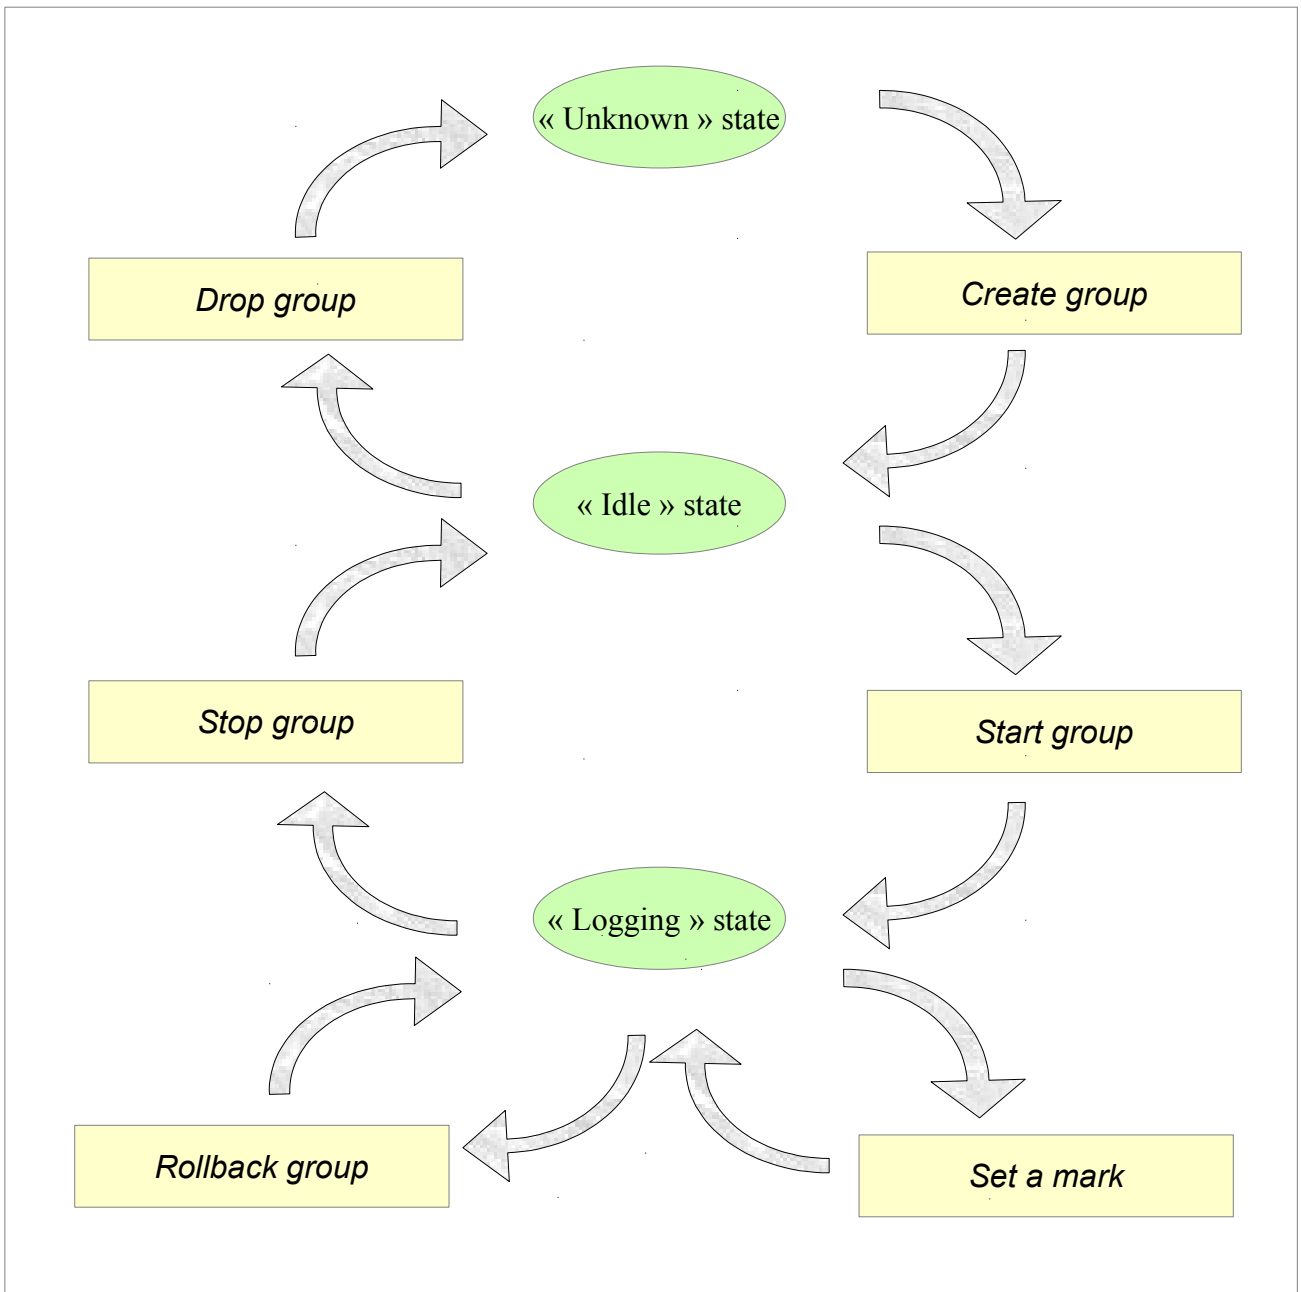
Except when explicitly noticed, the operations presented later can be indifferently executed by a superuser or by a role belonging to the *emaj\_adm* group.

## 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

### 4.2.2.1 The *emaj\_group\_def* table

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 has the following structure:

<b>Column</b>	<b>Type</b>	<b>Description</b>
<i>grpdef_group</i>	TEXT	tables group name
<i>grpdef_schema</i>	TEXT	name of the schema containing the application table or sequence
<i>grpdef_tblseq</i>	TEXT	application table or sequence name
<i>grpdef_priority</i>	INT	priority level for the table or sequence in E-Maj processing (optional)
<i>grpdef_log_dat_tsp</i>	TEXT	name of the tablespace containing the log table (optional)
<i>grpdef_log_idx_tsp</i>	TEXT	name of the tablespace containing the index of the log table (optional)

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

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.

### 4.2.2.2 Main columns

A tables group name (***grpdef\_group*** column) contains at least 1 character. It may contain spaces and/or any punctuation characters. But it is advisable to avoid commas, single or double quotes.

A table or a sequence of a given schema (***grpdef\_schema*** and ***grpdef\_tblseq*** columns) cannot be assigned to more than one 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 assigned to a group not created in “*audit\_only*” mode must have an explicit primary key (*PRIMARY KEY* clause in *CREATE TABLE* or *ALTER TABLE*).

E-Maj can process elementary partitions of partitionned tables created with the declarative DDL (with PostgreSQL 10+). They are processed as any other tables. However, as there is no need to protect mother tables, which remain empty, E-Maj refuses to include them in tables groups. All partitions of a partitionned table do not need to belong to a tables group. Partitions of a partitionned table can be assigned to different tables groups.

By their nature, *TEMPORARY TABLE* are not supported by E-Maj. *UNLOGGED* tables and tables created as *WITH OIDS* can only be members of “*audit\_only*” tables groups.

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 sequences should NOT be referenced in a tables group!

### 4.2.2.3 Optional columns

The type of the ***grpdef\_priority*** column is *INTEGER* and may be NULL. It defines a priority order in E-Maj tables processing. 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.

To optimize performances of E-Maj installations having a large number of tables, it may be useful to spread log tables and their index on several tablespaces. The ***grpdef\_log\_dat\_tsp*** column specifies the name of the tablespace to use for the log table of an application table. Similarly, the ***grpdef\_log\_idx\_tsp*** column specifies the name of the tablespace to use for the index of the log table.

If a column *grpdef\_log\_dat\_tsp* or *grpdef\_log\_idx\_tsp* is NULL (default value), the default tablespace of the current session at tables group creation is used.



For sequences, both *grpdef\_log\_dat\_tsp* and *grpdef\_log\_idx\_tsp* columns must be NULL.

### 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 abbreviated 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.

The function also creates the log schemas if needed.

On the contrary, if specific tablespaces are referenced for any log table or log index, these tablespaces must exist before the function's execution.

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 check the impact of this trigger on E-Maj rollbacks.

If a sequence of the group is associated either to a *SERIAL* or *BIGSERIAL* column or to a column created with a *GENERATED AS IDENTITY* clause, and the table that owns this column does not belong to the same tables group, the function also issues a *WARNING* message.

A specific version of the function allows to create an empty tables group, i.e. without any table or sequence at creation time:

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

The third parameter is false by default. If it is set to true, the group must not be referenced in the *emaj\_group\_def* table. Once created, an empty group can be then populated using the *emaj\_alter\_group()* function (see §4.3).

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 altered by the *emaj\_alter\_group()* function (see §4.3) or 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>',[<delete.old.logs?>]]);
```

The group must be first in *IDLE* state.

When a tables group is started, a first mark is created.

If specified, the initial 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 not specified, or is empty or *NULL*, a name is automatically generated: “START\_%”, where the '%' character represents the current transaction start time with a “hh.mn.ss.mmm” pattern.

The *<are.old.logs.to.be.deleted?>* parameter is an optional boolean. By default, its value is true, meaning that all log tables of the tables group are purged before the trigger activation. If the value is explicitly set to false, all rows from log tables are kept as is. The old marks are also preserved, even-though they are not usable for a rollback any more, (unlogged updates may have occurred while the tables group was stopped).

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 sets a *SHARE ROW 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.

The function also performs a purge of the oldest events in the *emaj\_hist* technical table (see §5.4).

When a group is started, its state becomes “*LOGGING*”.

Using the *emaj\_start\_groups()* function, several groups can be started at once:

```
SELECT emaj.emaj_start_groups('<group.names.array>',[  
'<mark.name>',[<delete.old.logs?>]]);
```

The chapter §4.9.3 explains how to describe the group names array.

#### 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 not specified or is empty or *NULL*, a name is automatically generated: “*MARK\_%*”, where the '%' character represents the current transaction start time with a “hh.mn.ss.mmm” pattern.

The function returns the number of tables and sequences contained in 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. The application sequences are processed first, to record their state as earlier as possible after the beginning of the transaction, these sequences not being protected against updates from concurrent transactions by any locking mechanism.

It is possible to set two consecutive marks without any update on any table between these marks.

The `emaj_set_mark_group()` function sets *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.

Using the `emaj_set_mark_groups()` function, a mark can be set on several groups at once:

```
SELECT emaj.emaj_set_mark_groups('<group.names.array>',  
'<mark.name>');
```

The chapter §4.9.3 explains how to describe the group names array.

## 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 * FROM emaj.emaj_rollback_group('<group.name>',  
'<mark.name>', <is_alter_group_allowed>);
```

The tables group must be in *LOGGING* state and the supplied mark must be usable for a rollback, i.e. it cannot be logically deleted.

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

The third parameter is a boolean that indicates whether the rollback operation may target a mark set before an alter group operation (function `emaj_alter_group()` - see §4.3). Depending on their nature, changes performed on tables groups in *LOGGING* state can be automatically cancelled or not. In some cases, this cancellation can be partial. By default, this parameter is set to *FALSE*.

The function returns a set of rows with a severity level set to either “*Notice*” or “*Warning*” values, and a textual message. The function returns a “*Notice*” row indicating the number of tables and sequences that have been **effectively** modified by the rollback operation. Other messages of type “*Warning*” may also be reported when the rollback operation has processed tables group changes.

To be sure that no concurrent transaction updates any table of the group during the rollback operation, the `emaj_rollback_group()` function explicitly sets an *EXCLUSIVE* lock on each table of the group. If transactions updating 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. But tables of the group remain accessible for read only transactions during the operation.

If tables belonging to the group to rollback have triggers, these triggers are by default temporarily disabled during the operation. Using the *emaj\_ignore\_app\_trigger()* function, it is possible to record triggers as 'not to be automatically disabled at rollback' (see §4.8.3 and §5.9.3).

If a table impacted by the rollback owns a *foreign key* or is referenced by a *foreign key* from another table, then this *foreign key* is taken into account by the rollback operation. If the check of the keys created or modified by the rollback cannot be deferred at the end of the operation (constraint not declared as *DEFERRABLE*), then this *foreign key* is dropped at the beginning of the rollback and recreated at the end.

When the volume of updates to cancel is high and the rollback operation is therefore long, it is possible to monitor the operation using the *emaj\_rollback\_activity()* function (§4.8.4.2) or the *emajRollbackMonitor.php* client (§4.11).

When the rollback operation is completed, the following are deleted:

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

The history of executed rollback operations is maintained into the *emaj\_rlbk* table. The final state of the operation is accessible from the *rlbk\_status* and *rlbk\_msg* columns of this *emaj\_rlbk* table.

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



By their nature, the reset of sequences 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.

Using the *emaj\_rollback\_groups()* function, several groups can be rolled back at once:

```
SELECT * FROM emaj.emaj_rollback_groups('<group.names.array>',  
'<mark.name>', <is_alter_group_allowed>);
```

The chapter §4.9.3 explains how to describe the group names array.

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.

An old version of these functions had only 2 input parameters and just returned an integer representing the number of effectively processed tables and sequences.

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

```
SELECT emaj.emaj_rollback_groups('<group.names.array>',  
'<mark.name>');
```

Both functions are deprecated and are subject to be deleted in a future E-Maj version.

#### 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 * FROM emaj.emaj_logged_rollback_group('<group.name>',  
'<mark.name>', <is_alter_group_allowed>);
```

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

The tables group must be in *LOGGING* state and the supplied mark must be usable for a rollback, i.e. it cannot be logically deleted.

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

The third parameter is a boolean that indicates whether the rollback operation may target a mark set before an alter group operation (function *emaj\_alter\_group()* - see §4.3). Depending on their nature, changes performed on tables groups in *LOGGING* state can be automatically cancelled or not. In some cases, this cancellation can be partial. By default, this parameter is set to *FALSE*.

The function returns a set of rows with a severity level set to either “*Notice*” or “*Warning*” values, and a textual message. The function returns a “*Notice*” row indicating the number of tables and sequences that have been **effectively** modified by the rollback operation. Other messages of type “*Warning*” may also be reported when the rollback operation has processed tables group changes.

To be sure that no concurrent transaction updates any table of the group during the rollback operation, the *emaj\_logged\_rollback\_group()* function explicitly sets an

*EXCLUSIVE* lock on each table of the group. If transactions updating 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. But tables of the group remain accessible for read only transactions during the operation.

If tables belonging to the group to rollback have triggers, these triggers are by default temporarily disabled during the operation. Using the *emaj\_ignore\_app\_trigger()* function, it is possible to record triggers as 'not to be automatically disabled at rollback' (see §4.8.3 and §5.9.3).

If a table impacted the rollback owns a *foreign key* or is referenced by a *foreign key* from another table, then this *foreign key* is taken into account by the rollback operation. If the check of the keys created or modified by the rollback cannot be deferred at the end of the operation (constraint not declared as *DEFERRABLE*), then this *foreign key* is dropped at the beginning of the rollback and recreated at the end.

Unlike with *emaj\_rollback\_group()* function, at the end of the operation, the log tables content as well as the marks following 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".

When the volume of updates to cancel is high and the rollback operation is therefore long, it is possible to monitor the operation using the *emaj\_rollback\_activity()* function (§4.8.4.2) or the *emajRollbackMonitor.php* client (§4.11).

The history of executed rollback operations is maintained into the *emaj\_rlbk* table. The final state of the operation is accessible from the *rlbk\_status* and *rlbk\_msg* columns of this *emaj\_rlbk* table.

Following the rollback operation, it is possible to resume updating the database, 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 delete 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)

A “*consolidation*” function for “*ogged rollback*” allows to transform a logged rollback into a simple unlogged rollback (see §4.4.6).

Using the *emaj\_logged\_rollback\_groups()* function, several groups can be rolled back at once:

```
SELECT * FROM emaj.emaj_logged_rollback_groups  
( '<group.names.array>', '<mark.name>', <is_alter_group_allowed> );
```

The chapter §4.9.3 explains how to describe the group names array.

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.

An old version of these functions had only 2 input parameters and just returned an integer representing the number of effectively processed tables and sequences.

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

```
SELECT emaj.emaj_logged_rollback_groups('<group.names.array>',  
'<mark.name>');
```

Both functions are deprecated and are subject to be deleted in a future E-Maj version.

## 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>[, '<mark.name>');
```

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



If the mark parameter is not specified or is empty or NULL, a mark name is generated: STOP\_% where % represents the current transaction start time expressed as “hh.mn.ss.mmm”.

Stopping a tables group simply deactivates log triggers of application tables of the group. The setting of *SHARE ROW EXCLUSIVE* locks 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 the *emaj\_stop\_group()* function for a tables group already stopped does not generate an error. Only a warning message is returned.

Using the *emaj\_stop\_groups()* function, several groups can be stopped at once:

```
SELECT emaj.emaj_stop_groups('<group.names.array>',[  
'<mark.name>']);
```

The chapter §4.9.3 explains how to describe the group names array.

## 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 first (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 in 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 function also drops all log schemas that are now useless.

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.

## 4.3 MODIFYING TABLES GROUPS

---

Several types of events may lead to alter a tables group:

- the tables group definition may change, some tables or sequences may have been added or suppressed,
- one of the E-Maj parameters linked to a table (priority, schema, tablespaces,...) may have been modified,
- the structure of one or several application tables of the tables group may have changed, such as an added or dropped column or a change in a column type.

### 4.3.1 Modifying a tables group in IDLE state

In all cases, the following steps can be performed:

- stop the group, if it is in *LOGGING* state, using the *emaj\_stop\_group()* function,
- update the *emaj\_group\_def* table and/or modify the application schema,
- drop and recreate the tables group, using the *emaj\_drop\_group()* and *emaj\_create\_group()* functions.

But this last step can be also performed by the *emaj\_alter\_group()* function, with a statement like:

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

The function returns the number of tables and sequences that now belong to the tables group.

The *emaj\_alter\_group()* function also recreates E-Maj objects that may be missing (log tables, functions, ...).

The function creates and drops the log schemas when needed.

Once altered, a tables group remains in IDLE state, but its log tables become empty.

The “*rollbackable*” or “*audit\_only*” characteristic of the tables group cannot be changed using the *emaj\_alter\_group()* function. To change it, the tables group must be dropped and re-created using the *emaj\_drop\_group()* and *emaj\_create\_group()* functions.

All actions that are chained by the *emaj\_alter\_group()* function are executed on behalf of a unique transaction. As a consequence, if an error occurs during the operation, the tables group remains in its previous state.

In most cases, executing the *emaj\_alter\_group()* function is much more efficient than chaining both *emaj\_drop\_group()* and *emaj\_create\_group()* functions.

It is possible to update the *emaj\_group\_def* table, when the tables group is in logging state. However it will not have an effect until the group is altered (or dropped and re-created).

Using the *emaj\_alter\_groups()* function, several groups can be modified at once:

```
SELECT emaj.emaj_alter_groups('<group.names.array>');
```

This function allows to move a table or a sequence from one tables group to another in a single operation.

The chapter §4.9.3 explains how to describe the group names array.

### 4.3.2 Modifying a tables group in LOGGING state

But the previous method has several drawbacks:

- logs recorded before the operation are lost,
- it is not possible to rollback a tables group to a previous state anymore.

However, some actions are possible while the tables groups are in *LOGGING* state. The following table lists these allowed actions. Details about each method is presented below.

Action	LOGGING group	Method
Change the log data tablespace	Yes	<i>emaj_group_def</i> update
Change the log index tablespace	Yes	<i>emaj_group_def</i> update
Change the E-Maj priority	Yes	<i>emaj_group_def</i> update
Remove a table/sequence from a group	Yes	<i>emaj_group_def</i> update
Add a table/sequence to a group	Yes	<i>emaj_group_def</i> update
Change the group ownership of a table/sequence	Yes	<i>emaj_group_def</i> update
Repair a table	Yes	Chaining remove/add
Rename a table	No	
Rename a sequence	No	
Change the schema of a table	No	
Change the schema of a sequence	No	
Rename a table's column	No	
Change a table's structure	No	
Other forms of ALTER TABLE	Yes	No E-Maj impact

### 4.3.2.1 The « emaj\_group\_def update » method

Most attributes of the *emaj\_group\_def* table describing the tables groups can be dynamically changed while groups have not been stopped.

To do this, the following steps can be performed:

- modify the *emaj\_group\_def* table,
- call one of the *emaj\_alter\_group()* or *emaj\_alter\_groups()* functions.

For tables groups in *LOGGING* state, these functions set a *ROW EXCLUSIVE* lock on each application table of these groups.

On these same tables groups, they also set a mark whose name can be supplied as parameter. The syntax of these calls becomes:

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

or

```
SELECT emaj.emaj_alter_groups('<group.names.array>' [, '<mark>');
```

If the parameter representing the mark is not specified, or is empty or *NULL*, a name is automatically generated: “ALTER\_%”, where the '%' character represents the current transaction start time with a “hh.mn.ss.mmm” pattern.

An E-Maj rollback operation targeting a mark set before such groups changes does NOT automatically cancel these changes.

However, the administrator can apply the same procedure to reset a tables group to a prior state.

### 4.3.2.2 Incidence of tables or sequences addition or removal in a group in *LOGGING* state



Once a table or a sequence is removed from a tables group, any rollback operation will leave this object unchanged.


Once unlinked from its tables group, the application table or sequence can be altered or dropped. The historical data linked to the object (logs, marks traces,...) are kept as is so that they can be later examined. However, they remain linked to the tables group that owned the object. To avoid any confusion, log tables are renamed, adding a numeric suffix to its name. These logs and marks traces will only be deleted by a group's reset operation (Cf §4.4.1) or by the deletion of the oldest marks of the group (Cf §4.5.5).



When a table or a sequence is added into a tables group in *LOGGING* state, it is then processed by any further rollback operation. But if the rollback operation targets a mark set before the addition into the group, the table or the sequence is left in its state at the time of the addition into the group and a warning message is issued. Such a table or sequence will not be processed by a SQL script generation function call if the requested start mark has been set before the addition of the table or sequence into the group

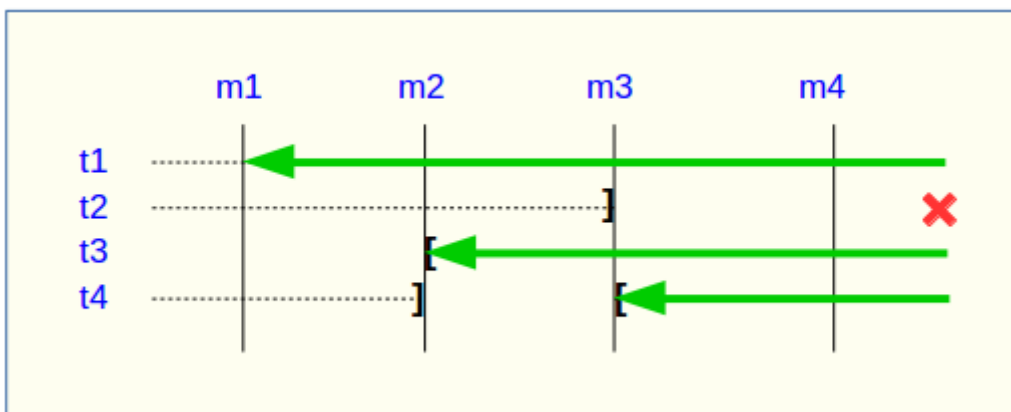
Some graphs help to more easily visualize the consequences of the addition or the removal of a table or a sequence into/from a tables group in *LOGGING* state.

Let's use a tables group containing 4 tables (t1 to t4) and 4 marks set over time (m1 to m4). At m2, t3 has been added to the group while t4 has been removed. At m3, t2 has been removed from the group while t4 has been re-added.



A rollback to the mark m1:

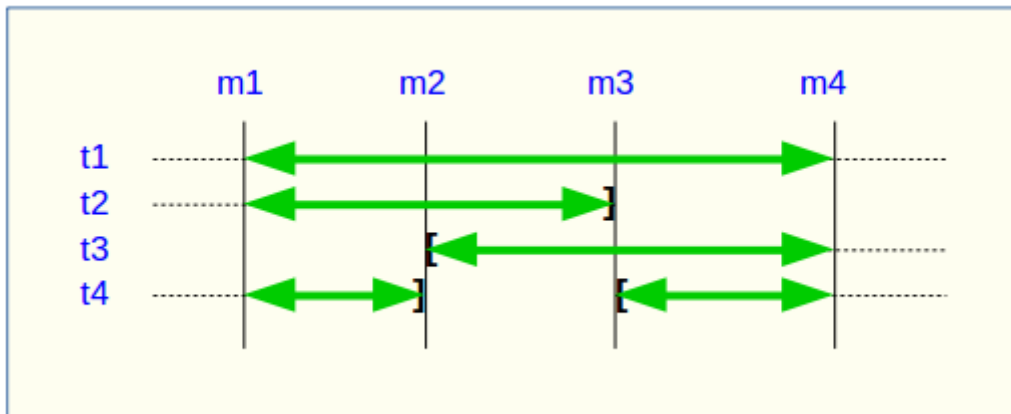
- would process the table t1,
- would **NOT** process the table t2, for lack of log after m3,
- would process the table t3, but only up to m2,
- would process the table t4, but only up to m3, for lack of log between m2 and m3.



A log statistics report between the marks m1 and m4 would contain:

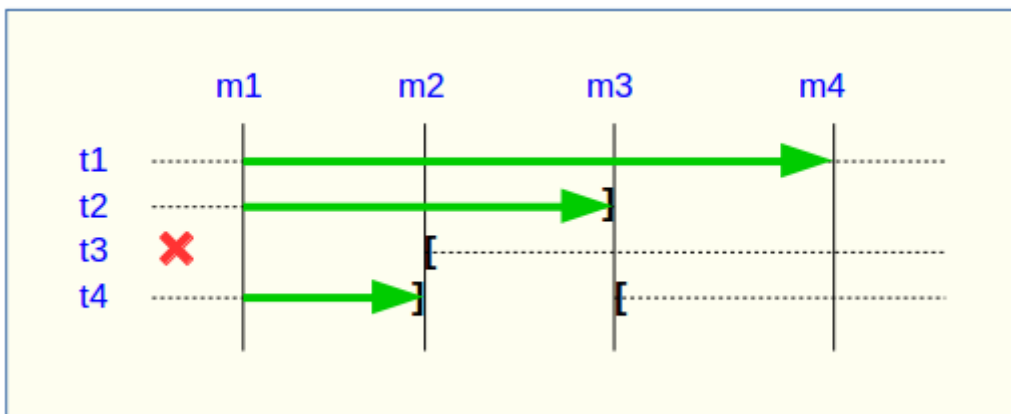
- 1 row for t1 (m1,m4),
- 1 row for t2 (m1,m3),

- 1 row for t3 (m2,m4),
- 2 rows for t4 (m1,m2) and (m3,m4).



The SQL script generation for the marks interval m1 to m4:

- would process the table t1,
- would process the table t2, but only up the mark m3,
- would **NOT** process the table t3, for lack of log before m2,
- would process the table t4, but only up to the mark m2, for lack of log between m2 and m3.



If the structure of an application table has been inadvertently changed while it belonged to a tables group in LOGGING state, the mark set and rollback operations will be blocked by the E-Maj internal checks. To avoid stopping, altering and then restarting the tables group, it is possible to only remove the concerned table from its group and then to re-add it.

When a table changes its affected group, the impact on the ability to generate a SQL script or to rollback the source and destination tables groups is similar to removing the table from its source group and then adding the table to the destination group.

#### 4.3.3 The "Chaining Remove/Add" method

Eventhough the event triggers created with E-Maj limit the risk, some E-Maj components that support an application table (log table, sequence or function) may be dropped. In such a case, the associated tables group cannot work correctly anymore.

In order to solve the issue without stopping the tables group (and thus loose the benefits of the recorded logs), it is possible to remove the table from its group and then re-add it. Four steps need to be executed:

- delete the row corresponding to the application table from the *emaj\_group\_def* table,
- call the *emaj\_alter\_group()* function for the related tables group, in order to effectively remove the table from the group,
- insert again the row corresponding to the table into the *emaj\_group\_def* table,
- call the *emaj\_alter\_group()* function again in order to re-add the table to the group.

Of course, once the table is removed from its group, the content of the associated logs cannot be used for a potential rollback or script generation anymore.

It may also happen that an application table or sequence be dropped. In this case, the table of sequence can be removed from its group, by chaining these steps:

- delete the row corresponding to the application table or sequence from the *emaj\_group\_def* table,
- call the *emaj\_alter\_group()* function for the related tables group, in order to effectively remove the table or sequence.



## 4.4 OTHER GROUPS MANAGEMENT FUNCTIONS

---

### 4.4.1 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.

### 4.4.2 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.

### 4.4.3 Protection of a tables group against rollbacks

It may be useful at certain time to protect tables groups against accidental rollbacks, in particular with production databases. Two functions fit this need.

The *emaj\_protect\_group()* function set a protection on a tables group.

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

The function returns the integer 1 if the tables group was not already protected, or 0 if it was already protected.

Once the group is protected, any logged or unlogged rollback attempt will be refused.

An “*audit\_only*” or “*idle*” tables group cannot be protected.

When a tables group is started, it is not protected. When a tables group that is protected against rollbacks is stopped, it loses its protection.

The *emaj\_unprotect\_group()* function remove an existing protection on a tables group.

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

The function returns the integer 1 if the tables group was previously protected, or 0 if it was not already protected.

An “*audit\_only*” tables group cannot be unprotected.

Once the protection of a tables group is removed, it becomes possible to execute any type of rollback operation on the group.

A protection mechanism at mark level complements this scheme (Cf §4.5.6).

#### 4.4.4 Forced stop of a tables group

It may occur that a corrupted tables group cannot be stopped. This may be the case for instance if an application table belonging to a tables group has been inadvertently dropped while the group was in *LOGGING* state. If usual *emaj\_stop\_group()* or *emaj\_stop\_groups()* functions return an error, it is possible to force a group stop using the *emaj\_force\_stop\_group()* function.

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

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

The *emaj\_force\_stop\_group()* function performs the same actions as the *emaj\_stop\_group()* function, except that:

- it supports the lack of table or E-Maj trigger to deactivate, generating a “*warning*” message in such a case,
- it does NOT set a stop mark.

Once the function is completed, the tables group is in *IDLE* state. It may then be altered or dropped, using the *emaj\_alter\_group()* or *emaj\_drop\_group()* functions.

It is recommended to only use this function if it is really needed.

#### 4.4.5 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.

Note: Since the *emaj\_force\_stop\_group()* function has been created, this *emaj\_force\_drop\_group()* function becomes useless. It may be removed in a future version.

#### 4.4.6 Logged rollback consolidation

Following the execution of a “*logged rollback*”, and once the rollback operation recording becomes useless, it is possible to “*consolidate*” this *rollback*, meaning to some extent to transform it into “*unlogged rollback*”. At the end of the consolidation operation, marks and logs between the rollback target mark and the end rollback mark are deleted. The *emaj\_consolidate\_rollback\_group()* function fits this need.

```
SELECT emaj.emaj_consolidate_rollback_group('<group.name>',  
<end.rollback.mark>);
```

The concerned logged rollback operation is identified by the name of the mark generated at the end of the rollback. This mark must always exist, but may have been renamed.

The '*EMAJ\_LAST\_MARK*' keyword may be used as mark name to reference the last set mark.

The *emaj\_get\_consolidable\_rollbacks()* function may help to identify the rollbacks that may be consolidated (See §4.4.7).

Like rollback functions, the *emaj\_consolidate\_rollback\_group()* function returns the number of effectively processed tables and sequences.

The tables group may be in *LOGGING* or *IDLE* state.

The rollback target mark must always exist but may have been renamed. However, intermediate marks may have been deleted.

When the consolidation is complete, only the rollback target mark and the end rollback mark are kept.

The disk space of deleted rows will become reusable as soon as these log tables will be “vacuumed”.

Of course, once consolidated, a “*logged rollback*” cannot be cancelled (or rolled back) any more, the start rollback mark and the logs covering this rollback being deleted.

The consolidation operation is not sensitive to the protections set on groups or marks, if any.

If a database has enough disk space, it may be interesting to replace a simple *unlogged rollback* by a *logged rollback* followed by a *consolidation* so that the application tables remain readable during the rollback operation, thanks to the lower locking mode used for *logged rollbacks*.

#### 4.4.7 List of “*consolidable rollbacks*”

The `emaj_get_consolidable_rollbacks()` function help to identify the rollbacks that may be consolidated.

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

The function returns a set of rows with the following columns:

- |  |   |
|--|---|
| ➤ <code>cons_group</code>                    | rolled back tables group                  |
| ➤ <code>cons_target_rlbk_mark_name</code>    | rollback target mark name                 |
| ➤ <code>cons_target_rlbk_mark_time_id</code> | temporal reference of the target mark (*) |
| ➤ <code>cons_end_rlbk_mark_name</code>       | rollback end mark name                    |
| ➤ <code>cons_end_rlbk_mark_time_id</code>    | temporal reference of the end mark (*)    |
| ➤ <code>cons_rows</code>                     | number of intermediate updates            |
| ➤ <code>cons_marks</code>                    | number of intermediate marks              |

(\*) `emaj_time_stamp` table identifiers ; this table contains the time stamps of the most important events of the tables groups life.

Using this function, it is easy to consolidate at once all “consolidable” rollbacks for all tables groups in order to recover as much as possible disk space:

```
SELECT emaj.emaj_consolidate_rollback_group(cons_group,  
cons_end_rlbk_mark_name) FROM  
emaj.emaj_get_consolidable_rollbacks();
```

The *emaj\_get\_consolidable\_rollbacks()* function may be used by *emaj\_adm* and *emaj\_viewer* roles.

## 4.5 MARKS MANAGEMENT FUNCTIONS

---

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

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 web clients (See §6). Indeed, they systematically display the comments in the groups marks list.

### 4.5.2 Search a mark

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

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

or

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

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

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

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

### 4.5.3 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.

### 4.5.4 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.

If a table has been detached from a tables group (Cf. 4.3.2), and the deleted mark corresponds to the last known mark for this table, the logs for the period between this mark and the preceding one are deleted,

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

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.

The function returns the number of deleted marks.

The function also performs a purge of the oldest events in the *emaj\_hist* technical table (see §5.4).

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

However, as the log rows deletion cannot use any *TRUNCATE* command (unlike with the *emaj\_start\_group()* or *emaj\_reset\_group()* functions), using *emaj\_delete\_before\_mark\_group()* function may take a longer time than simply stopping and restarting the group. In return, no lock is set on the tables of the group. Its execution may continue while other processes update the application tables. Nothing but other E-Maj operations on the same tables group, like setting a new mark, would wait until the end of an *emaj\_delete\_before\_mark\_group()* function execution.

When associated, the functions *emaj\_delete\_before\_mark\_group()* and *emaj\_get\_previous\_mark\_group()* allow to delete marks older than a retention delay. 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));
```

## 4.5.6 Protection of a mark against rollbacks

To complement the mechanism of tables group protection against accidental rollbacks (see §4.4.3), it is possible to set protection at mark level. Two functions fit this need.

The *emaj\_protect\_mark\_group()* function sets a protection on a mark for a tables group.

```
SELECT  
emaj.emaj_protect_mark_group('<groupe.name>', '<mark.name>');
```

The function returns the integer 1 if the mark was not previously protected, or 0 if it was already protected.

Once a mark is protected, any logged or unlogged rollback attempt is refused if it reset the tables group in a state prior this protected mark.

A mark of an “audit-only” or an “idle” tables group cannot be protected.

When a mark is set, it is not protected. Protected marks of a tables group automatically loose their protection when the group is stopped. Warning: deleting a protected mark also deletes its protection. This protection is not moved on an adjacent mark.

The *emaj\_unprotect\_mark\_group()* function remove an existing protection on a tables group mark.

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

The function returns the integer 1 if the mark was previously protected, or 0 if it was not yet protected.

A mark of an “audit-only” tables group cannot be unprotected.

Once a mark protection is removed, it becomes possible to execute any type of rollback on a previous mark.



## 4.6 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()* provides more detailed information than *emaj\_log\_stat\_group()*, the number of updates been reported per table, SQL type (INSERT/UPDATE/DELETE) and connection role.

Two other E-Maj functions, *emaj\_estimate\_rollback\_group()* and *emaj\_estimate\_rollback\_groups()*, provide an estimate of how long a rollback for one or several groups to a given mark may last.

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

### 4.6.1 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>', '<end.mark>');
```

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\_first\_mark*: mark name of the period start (type TEXT),
- *stat\_first\_mark\_datetime*: mark timestamp of the period start (type TIMESTAMPTZ),
- *stat\_last\_mark*: mark name of the period end (type TEXT),
- *stat\_last\_mark\_datetime*: mark timestamp of the period end (type TIMESTAMPTZ),
- *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.

Most of the time, the *stat\_first\_mark*, *stat\_first\_mark\_datetime*, *stat\_last\_mark* and *stat\_last\_mark\_datetime* columns reference the start and end marks of the requested period. But they can contain other values when a table has been added or removed from the tables group during the requested time interval.

It is possible to easily execute more precise requests on these statistics. For instance, 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 table updates before being cancelled.

Using the *emaj\_log\_stat\_groups()* function, log statistics can be obtained for several groups at once:

```
SELECT emaj.emaj_log_stat_groups('<group.names.array>',
'<start.mark>', '<end.mark>');
```

The chapter §4.9.3 explains how to describe the group names array.

## 4.6.2 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>', '<end.mark>');
```

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\_first\_mark* : mark name of the period start (type TEXT),
- *stat\_first\_mark\_datetime* : mark timestamp of the period start (type TIMESTAMPTZ),
- *stat\_last\_mark* : mark name of the period end (type TEXT),
- *stat\_last\_mark\_datetime* : mark timestamp of the period end (type TIMESTAMPTZ),
- *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()*, the *emaj\_detailed\_log\_stat\_group()* function doesn't return any rows for tables having no logged updates inside the requested marks range. So *stat\_rows* column never contains 0.

Most of the time, the *stat\_first\_mark*, *stat\_first\_mark\_datetime*, *stat\_last\_mark* and *stat\_last\_mark\_datetime* columns reference the start and end marks of the requested period. But they can contain other values when a table has been added or removed from the tables group during the requested time interval.

Using the *emaj\_detailed\_log\_stat\_groups()* function, detailed log statistics can be obtained for several groups at once:

```
SELECT emaj.emaj_detailed_log_stat_groups('<group.names.array>',
'<start.mark>', '<end.mark>');
```

The chapter §4.9.3 explains how to describe the group names array.

### 4.6.3 Estimate the rollback duration

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

```
SELECT emaj.emaj_estimate_rollback_group('<group.name>',
'<mark.name>', <is.logged>);
```

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

The third parameter indicates whether the E-Maj rollback to simulate is a logged rollback or not.

The function returns an *INTERVAL* value.

The tables group must be in *LOGGING* state and the supplied mark must be usable for a rollback, i.e. it cannot be logically deleted.

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,
- 6 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.

Using the *emaj\_estimate\_rollback\_groups()* function, it is possible to estimate the duration of a rollback operation on several groups:

```
SELECT emaj.emaj_estimate_rollback_groups('<group.names.array>',  
'<mark.name>', <is.logged>);
```

The chapter §4.9.3 explains how to describe the group names array.

## 4.7 DATA EXTRACTION FUNCTIONS

---

Three functions extract data from E-Maj infrastructure and store them into external files.

### 4.7.1 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 to files all tables and sequences of a group with:

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

The directory/folder name must be supplied as an absolute pathname and must have been previously created. This directory/folder must have the appropriate permission so that the PostgreSQL instance can write in it.

The third parameter defines the output files format. It is a character string that matches the precise syntax available for the COPY TO SQL statement.

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.

New files will overwrite existing files of the same name.

Created files are named with the following pattern:

*<schema.name>\_<table/sequence.name>.snap*

Some inconvenient in file name characters, namely spaces, “/”, “\”, “\$”, “>”, “<”, and “\*” are replaced by “\_”.

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

Files corresponding to tables contain one record per row, in the format corresponding to the supplied parameter. These records are sorted in ascending order of the primary key.

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.

## 4.7.2 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 provides a way to archive updates performed by one or more previous operations. 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>', '<COPY.options>');
```

A *NULL* value or an empty string may be used as start mark, representing the first known mark.

A *NULL* value or an empty string may be used as end mark, representing the current situation.

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

The directory/folder name must be supplied as an absolute pathname and must have been previously created. This directory/folder must have the appropriate permission so that the PostgreSQL instance can write in it.

The fifth parameter defines the output files format. It is a character string that matches the precise syntax available for the COPY TO SQL statement.

The function returns the number of generated files.

This *emaj\_snap\_log\_group()* function generates one file per log table, containing the part of this table that corresponds to the updates performed between both supplied marks. Created files name has the following pattern:

*<log.table.name>.snap*

So most of the time, they look like:

*<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>`

All these files are stored in the directory or folder corresponding to the fourth parameter. New files will overwrite existing files of the same name.

Some inconvenient in file name characters, namely spaces, “/”, “\”, “\$”, “>”, “<”, and “\*” are replaced by “\_”.

At the end of the operation, a file named `_INFO` is created in this same directory/folder. It contains a message including the table's group name, the mark's 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. If no end mark has been supplied, the log tables snap is bounded by a pseudo mark set at the function start. This ensures that, if the group is in logging state, output files will not contain updates recorded after the function start.

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 described in §5.2.1.

### 4.7.3 SQL script generation to replay logged updates

Log tables contain all needed information to replay updates. Therefore, it is possible to generate SQL statements corresponding to all updates that occurred between two marks or between a mark and the current situation, and record them into a file. This is the purpose of the `emaj_gen_sql_group()` function.

So these updates can be replayed after the corresponding tables have been restored in their state at the initial mark, without being obliged to rerun application programs.

To generate this SQL script, just execute the following statement:

```
SELECT emaj.emaj_gen_sql_group('<group.name>', '<start.mark>',  
'<end.mark>', '<file>' [, <tables/sequences.array>);
```

A `NULL` value or an empty string may be used as start mark, representing the first known mark.

A `NULL` value or an empty string may be used as end mark, representing the current situation.

The keyword `'EMAJ_LAST_MARK'` can be used as mark name, representing the last set mark.

The output file name must be supplied as an absolute pathname. It must have the appropriate permission so that the PostgreSQL instance can write to it. If the file already exists, its content is overwritten.

The last parameter is optional. It allows filtering of the tables and sequences to process. If the parameter is omitted or has a *NULL* value, all tables and sequences of the tables group are processed. If specified, the parameter must be expressed as a non empty array of text elements, each of them representing a schema qualified table or sequence name. Both syntaxes can be used:

```
ARRAY['sch1.tbl1','sch1.tbl2']
```

or

```
{ "sch1.tbl1" , "sch1.tbl2" }
```

The function returns the number of generated statements (not including comments and transaction management statements).

The tables group may be in *IDLE* state while the function is called.

In order to generate the script, all tables must have an explicit *PRIMARY KEY*.



If a tables and sequences list is specified to limit the *emaj\_gen\_sql\_group()* function's work, it is the user's responsibility to take into account the possible presence of foreign keys, in order to let the function produce a viable SQL script.

All statements, *INSERT*, *UPDATE*, *DELETE* and *TRUNCATE* (for *audit\_only* tables groups), are generated in the order of their initial execution.

The statements are inserted into a single transaction. They are surrounded by a *BEGIN TRANSACTION;* statement and a *COMMIT;* statement. An initial comment specifies the characteristics of the script generation: generation date and time, related tables group and used marks.

*TRUNCATE* statements recorded for *audit\_only* tables groups are also included into the script.

At the end of the script, sequences belonging to the tables group are set to their final state.

Then, the generated file may be executed as is by *psql* tool, using a connection role that has enough rights on accessed tables and sequences.

The used technology may result to doubled backslashes in the output file. These doubled characters must be suppressed before executing the script, for instance, in Unix/Linux environment, using a command like:

```
sed 's/\\\\\\\\/g' <file.name> | psql ...
```

As the function can generate a large or even very large file (depending on the log volume), it is the user's responsibility to provide a sufficient disk space.



It is also the user's responsibility to deactivate application triggers, if any exist, before executing the generated script.

Using the *emaj\_gen\_sql\_groups()* function, it is possible to generate a sql script related to several groups:

```
SELECT emaj.emaj_gen_sql_groups('<group.names.array>',  
'<start.mark>', '<end.mark>', '<file>' [, <tables/sequences.array>);
```

The chapter §4.9.3 explains how to describe the group names array.

## 4.8 OTHER FUNCTIONS

---

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

A function is also available to check the consistency of the E-Maj environment. It consists in checking the integrity of all E-Maj schemas and all created tables groups. This function can be called with the following SQL statement:

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

For each E-Maj schema (*emaj* and each log schema) the function verifies that:

- all tables, functions, sequences and types contained in the schema are either objects of the extension, or linked to created tables groups,
- they don't contain any view, foreign table, domain, conversion, operator or operator class.

Then, for each created tables group, the function performs the same checks as those performed when a group is started, a mark is set, or a rollback is executed (see §5.3.1).

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

*'No error detected'*

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.

### 4.8.2 Getting the current log table linked to an application table

The *emaj\_get\_current\_log\_table()* function allows to get the schema and table names of the current log table linked to a given application table.

```
SELECT log_schema, log_table FROM  
emaj_get_current_log_table(<schema>, <table>);
```

The function always returns 1 row. If the application table does not currently belong to any tables group, the *log\_schema* and *log\_table* columns are set to NULL.

The *emaj\_get\_current\_log\_table()* function can be used by *emaj\_adm* and *emaj\_viewer* E-Maj roles.

It is possible to build a statement accessing a log table. For instance:

```
SELECT 'select count(*) from '  
      || quote_ident(log_schema) || '.' || quote_ident(log_table)  
      FROM emaj.emaj_get_current_log_table('myschema','mytable');
```

### 4.8.3 Not disabling application triggers at E-Maj rollback time

Application triggers are automatically disabled during E-Maj rollback operations. Under some circumstances, it may be desirable to keep them enabled (See §5.9.3 for more details). The *emaj\_ignore\_app\_trigger()* function achieves this. It allows to add or remove triggers into/from a list of triggers that do not need to be disabled during rollback operations.

```
SELECT emaj.emaj_ignore_app_trigger(<action>, <schema.name>,  
<table.name>, <trigger.name>);
```

The <action> parameter accepts 2 values: 'ADD' to add a trigger to the list or 'REMOVE' to delete a trigger from the list.

The trigger is identified by the 3 components: schema name, table name and trigger name.

The trigger name may contain '%' and '\_' wildcard characters. These characters have the same meaning as in the *LIKE* clause of the SQL language. Thus several triggers of a single table can be processed by a unique function call.

The function returns the number of triggers effectively added or removed.

The function does not process E-Maj triggers (log triggers or triggers protecting against *TRUNCATE*).

The triggers referenced as "not to be automatically disabled during E-Maj rollbacks" are registered into the *emaj.emaj\_ignored\_app\_trigger* table. This table contains 3 columns:

- *trg\_schema*            schema of the table holding the trigger
- *trg\_table*            table holding the trigger
- *trg\_name*             trigger name

In order to know the list of registered triggers, just display the table's content.

### 4.8.4 Monitoring rollback operations

When the volume of recorded updates to cancel leads to a long rollback, it may be interesting to monitor the operation to appreciate how it progresses. A function, named

*emaj\_rollback\_activity()*, and a client, *emajRollbackMonitor.php* (see §4.11), fit this need.

#### 4.8.4.1 Prerequisite

To allow E-Maj administrators to monitor the progress of a rollback operation, the activated functions update several technical tables as the process progresses. To ensure that these updates are visible while the transaction managing the rollback is in progress, they are performed through a *dblink* connection.

As a result, monitoring rollback operations requires the installation of the *dblink* extension (§3.3.2) as well as the insertion of a connection identifier usable by *dblink* into the *emaj\_param* table.

Recording the connection identifier can be performed with a statement like:

```
INSERT INTO emaj.emaj_param (param_key, param_value_text)
VALUES ('dblink_user_password','user=<user> password=<password>');
```

The declared connection role must have been granted the *emaj\_adm* rights (or be a superuser).

Lastly, the main transaction managing the rollback operation must be in a “read committed” concurrency mode (the default value).

#### 4.8.4.2 Monitoring function

The *emaj\_rollback\_activity()* function allows to see the progress of rollback operations.

Invoke it with the following statement:

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

The function does not require any input parameter.

It returns a set of rows of type *emaj.emaj\_rollback\_activity\_type*. Each row represents an in progress rollback operation, with the following columns:

- *rlbk\_id*                      rollback identifier
- *rlbk\_groups*                tables groups array associated to the rollback
- *rlbk\_mark*                    mark to rollback to
- *rlbk\_mark\_datetime*        date and time when the mark to rollback to has been set
- *rlbk\_is\_logged*              boolean taking the “true” value for logged rollbacks
- *rlbk\_nb\_session*            number of parallel sessions
- *rlbk\_nb\_table*                number of tables contained in the processed tables groups
- *rlbk\_nb\_sequence*        number of sequences contained in the processed tables groups
- *rlbk\_eff\_nb\_table*        number of tables having updates to cancel

- *rlbk\_status*                rollback operation state
- *rlbk\_start\_date*            rollback operation start timestamp
- *rlbk\_elapse*                elapse time spent since the rollback operation start
- *rlbk\_remaining*            estimated remaining duration
- *rlbk\_completion\_pct*        estimated percentage of the completed work

An in progress rollback operation is in one of the following state:

- *PLANNING*                the operation is in its initial planning phase,
- *LOCKING*                 the operation is setting locks,
- *EXECUTING*               the operation is currently executing one of the planned steps.

If the functions executing rollback operations cannot use *dblink* connections (extension not installed, missing or incorrect connection parameters,...), the *emaj\_rollback\_activity()* does not return any rows.

The remaining duration estimate is approximate. Its precision is similar to the precision of the *emaj\_estimate\_rollback\_group()* function (§4.6.3).

#### 4.8.5 Updating rollback operations state

The *emaj\_rlbk* technical table and its derived tables contain the history of E-Maj rollback operations.

When rollback functions cannot use *dblink* connections (see the conditions at §4.8.4.1), all updates of these technical tables are all performed inside a single transaction. Therefore:

- any rollback operation that has not been completed is invisible in these technical tables,
- any rollback operation that has been validated is visible in these technical tables with a “*COMMITTED*” state.

When rollback functions can use *dblink* connections, all updates of *emaj\_rlbk* and its related tables are performed in autonomous transactions. In this working mode, rollback functions leave the operation in a “*COMPLETED*” state when finished. A dedicated internal function is in charge of transforming the “*COMPLETED*” operations either into a “*COMMITTED*” state or into an “*ABORTED*” state, depending on how the main rollback transaction has ended. This function is automatically called when a new mark is set and when the rollback monitoring function is used.

If the E-Maj administrator wishes to check the status of recently executed rollback operations, he can use the *emaj\_cleanup\_rollback\_state()* function at any time:

```
SELECT emaj.emaj_cleanup_rollback_state();
```

The function returns the number of modified rollback operations.

## 4.8.6 Deactivating or reactivating event triggers

The E-Maj extension installation procedure activates event triggers to protect it (See §5.3.2). Normally, these triggers must remain in their state. But if the E-Maj administrator needs to deactivate and the reactivate them, he can use 2 dedicated functions.

To deactivate the existing event triggers:

```
SELECT emaj.emaj_disable_protection_by_event_triggers();
```

The function returns the number of deactivated event triggers (this value depends on the installed PostgreSQL version).

To reactivate existing event triggers:

```
SELECT emaj.emaj_enable_protection_by_event_triggers();
```

The function returns the number of reactivated event triggers.

## 4.9 MULTI-GROUPS FUNCTIONS

### 4.9.1 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.

### 4.9.2 Functions list

The following table lists the multi-groups functions, with their relative mono-group functions, some of them being discussed later.

Multi-groups functions	Relative mono-group function	§
<i>emaj.emaj_alter_groups()</i>	<i>emaj.emaj_alter_group()</i>	4.3
<i>emaj.emaj_start_groups()</i>	<i>emaj.emaj_start_group()</i>	4.2.4
<i>emaj.emaj_stop_groups()</i>	<i>emaj.emaj_stop_group()</i>	4.2.8
<i>emaj.emaj_set_mark_groups()</i>	<i>emaj.emaj_set_mark_group()</i>	4.2.5
<i>emaj.emaj_rollback_groups()</i>	<i>emaj.emaj_rollback_group()</i>	4.2.6
<i>emaj.emaj_logged_rollback_groups()</i>	<i>emaj.emaj_logged_rollback_group()</i>	4.2.7
<i>emaj.emaj_estimate_rollback_groups()</i>	<i>emaj.emaj_estimate_rollback_group()</i>	4.6.3
<i>emaj.emaj_log_stat_groups()</i>	<i>emaj.emaj_log_stat_group()</i>	4.6.1
<i>emaj.emaj_detailed_log_stat_groups()</i>	<i>emaj.emaj_detailed_log_stat_group()</i>	4.6.2
<i>emaj.emaj_gen_sql_groups()</i>	<i>emaj.emaj_gen_sql_group()</i>	4.7.3

The parameters of multi-groups functions are the same as those of their related mono-group function, except the first one. The *TEXT* table group parameter is replaced by a *TEXT ARRAY* parameter representing a tables groups list.

### 4.9.3 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 single 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.

#### 4.9.4 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. This may allows a set oriented build of this list, using for instance the *array\_agg()* function.

A tables groups list may contain duplicate values, NULL values or empty strings. These NULL values or empty strings are simply ignored. If a tables group name is listed several times, only one occurrence is kept.

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.



## 4.10 PARALLEL ROLLBACK CLIENT

---

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

### 4.10.1 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 its own thread.

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

To build the 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.

### 4.10.2 Prerequisites

Two equivalent tools are actually provided, one coded in php and the other in perl. Both need that some software components be installed on the server that executes the command (which is not necessarily the same as the one that hosts the PostgreSQL instance):

- for the php client, the *php* software and its PostgreSQL interface,
- for the perl client, the perl software with the *DBI* and *DBD::Pg* modules.

Rolling back 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.

### 4.10.3 Syntax

Both php and perl commands share the same syntax:

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

and

```
emajParallelRollback.pl -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)
- a specifies that the requested rollback is allowed to reach a mark set before an alter group operation (see §4.3)
- 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 allow the rollback operation to 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.

It is possible to monitor the multi-session rollback operations with the same tools as for mono-session rollbacks.

In order to test both *emajParallelRollback* commands, the E-Maj extension supplies a test script, *emaj\_prepare\_parallel\_rollback\_test.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.

#### 4.10.4 Examples

The command:

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

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

The command:

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

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

## 4.11 ROLLBACK MONITORING CLIENT

---

E-Maj delivers an external client to run as a command that monitors the progress of rollback operations in execution.

### 4.11.1 Prerequisite

Two equivalent tools are actually provided, one coded in php and the other in perl. Both need that some software components be installed on the server that executes the command (which is not necessarily the same as the one that hosts the PostgreSQL instance):

- for the php client, the *php* software and its PostgreSQL interface,
- for the perl client, the perl software with the *DBI* and *DBD::Pg* modules.

### 4.11.2 Syntax

Both php and perl commands share the same syntax:

```
emajRollbackMonitor.php [OPTIONS]...  
and  
emajRollbackMonitor.pl [OPTIONS]...
```

General options:

- i time interval between 2 displays (in seconds, default = 5s)
- n number of displays (default = 1)
- a maximum time interval for rollback operations to display (in hours, default = 24h)
- l maximum number of completed rollback operations to display (default = 3)
- 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.

The supplied connection role must either be a super-user or have *emaj\_adm* or *emaj\_viewer* rights.

For security reasons, it is not recommended to use the `-W` option to supply a password. Rather, it is advisable to use the `.pgpass` file (see PostgreSQL documentation).

## Examples

The command:

```
./client/emajRollbackMonitor.php -i 3 -n 10
```

displays 10 times and every 3 seconds, the list of in progress rollback operations and the list of the at most 3 latest rollback operations completed in the latest 24 hours.

The command:

```
./client/emajRollbackMonitor.pl -a 12 -l 10
```

displays only once the list of in progress rollback operations and the list of at most 10 operations completed in the latest 12 hours.

Example of display:

```
E-Maj (version 2.2.0) - Monitoring rollbacks activity
-----
04/02/2018 - 12:07:17
** rollback 34 started at 2018-02-04 12:06:20.350962+02 for groups {myGroup1,myGroup2}
   status: COMMITTED ; ended at 2018-02-04 12:06:21.149111+02
** rollback 35 started at 2018-02-04 12:06:21.474217+02 for groups {myGroup1}
   status: COMMITTED ; ended at 2018-02-04 12:06:21.787615+02
-> rollback 36 started at 2018-02-04 12:04:31.769992+02 for groups {group1232}
   status: EXECUTING ; completion 89 % ; 00:00:20 remaining
-> rollback 37 started at 2018-02-04 12:04:21.894546+02 for groups {group1233}
   status: LOCKING ; completion 0 % ; 00:22:20 remaining
-> rollback 38 started at 2018-02-04 12:05:21.900311+02 for groups {group1234}
   status: PLANNING ; completion 0 %
```

## 5 MISCELLANEOUS

---

### 5.1 PARAMETERS

---

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

The *emaj\_param* table structure is the following:

Column	Type	Description
<i>param_key</i>	TEXT	keyword identifying the parameter
<i>param_value_text</i>	TEXT	parameter value, if its type is text (otherwise NULL)
<i>param_value_numeric</i>	NUMERIC	parameter value, if its type is numeric (otherwise NULL)
<i>param_value_boolean</i>	BOOLEAN	parameter value, if its type is boolean (otherwise NULL)
<i>param_value_interval</i>	INTERVAL	parameter value, if its type is time interval (otherwise NULL)

The E-Maj extension installation procedure inserts a single row into the *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 the *emaj\_param* table to change the default value of some parameters.

Presented in alphabetic order, the existing key values are:

- *alter\_log\_table* (text) *ALTER TABLE* directive executed at the log table creation ; no *ALTER TABLE* executed by default (see §5.2.2).
- *avg\_fkey\_check\_duration* (interval) default value = 20  $\mu$ s ; defines the average duration of a *foreign key* value check ; can be modified to better represent the performance of the server that hosts the database (see §4.6.3).
- *avg\_row\_delete\_log\_duration* (interval) default value = 10  $\mu$ s ; 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 §4.6.3).
- *avg\_row\_rollback\_duration* (interval) default value = 100  $\mu$ s ; defines the average duration of a row rollback ; can be modified to better represent the performance of the server that hosts the database (see §4.6.3).
- *fixed\_dblink\_rollback\_duration* (interval) default value = 4 ms ; defines an additional cost for each rollback step when a *dblink* connection is used ; can be modified to better represent the performance of the server that hosts the database (see §4.6.3).

- *fixed\_table\_rollback\_duration* (interval) default value = 1 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 §4.6.3).
- *fixed\_step\_rollback\_duration* (interval) default value = 2,5 ms ; defines a fixed cost for each rollback step ; can be modified to better represent the performance of the server that hosts the database (see §4.6.3).
- *history\_retention* (interval) default value = 1 year ; it can be adjusted to change the retention delay of rows in the *emaj\_hist* history table (see § 5.4),

Below is an example of a SQL statement that defines a retention delay of history table's rows equal to 3 months:

```
INSERT INTO emaj.emaj_param (param_key, param_value_interval)
VALUES ('history_retention','3 months'::interval);
```

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

Only super-user and roles having *emaj\_adm* rights can access the *emaj\_param* table.

Roles having *emaj\_viewer* rights can only access a part of the *emaj\_param* table, through the *emaj.emaj\_visible\_param* view. This view just masks the real value of the *param\_value\_text* column for the '*dblink\_user\_password*' key.

## 5.2 LOG TABLES STRUCTURE

---

### 5.2.1 Standart structure

The structure of log tables is directly derived from the structure of the related application tables. The log tables 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\_gid* log row identifier
- *emaj\_changed* log row insertion timestamp
- *emaj\_txid* transaction id (the PostgreSQL \*txid\*) that performed the update
- *emaj\_user* connection role that performed the update

### 5.2.2 Adding technical columns

It is possible to add one or several technical columns to enrich the traces. These columns value must be set as a default value (a *DEFAULT* clause) associated to a function (so that the log triggers are not impacted).

To add one or several technical columns, a parameter of key '*alter\_log\_table*' must be inserted into the *emaj\_param* table (Cf §5.1). The associated text value must contain an ALTER TABLE clause. At the log table creation time, if the parameter exists, an ALTER TABLE statement with this parameter is executed.

For instance, one can add to log tables a column to record the value of the '*application\_name*' connection field with:

```
INSERT INTO emaj.emaj_param (param_key, param_value_text) VALUES
('alter_log_table', 'ADD COLUMN extra_col_appname TEXT
DEFAULT current_setting("application_name");');
```

Several ADD COLUMN directives may be concatenated, separated by a comma. For instance, to create columns recording the ip adress and port of the connected client:

```
INSERT INTO emaj.emaj_param (param_key, param_value_text) VALUES
('alter_log_table', 'ADD COLUMN emaj_user_ip INET DEFAULT inet_client_addr(),
ADD COLUMN emaj_user_port INT DEFAULT inet_client_port());');
```

To change the structure of existing log tables once the *alter\_log\_table* parameter has been set, the tables groups must be dropped and then recreated.

## 5.3 RELIABILITY

---

Two additional elements help in ensuring the E-Maj reliability: internal checks are performed at some key moments of tables groups life and event triggers can block some risky operations.

### 5.3.1 Internal checks

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

These tables group integrity checks verify that:

- the PostgreSQL version at tables group creation time is compatible with the current version,
- each application sequence or table of the group always exists,
- each table of the group has its log table, its log function and its triggers,



- the log tables structure always reflects the related application tables structure, and contains all required technical columns,
- for *ROLLBACKABLE* tables groups, no table has been altered as *UNLOGGED* or *WITH OIDS*,
- for *ROLLBACKABLE* tables groups, application tables have their primary key and their structure has not changed.

By using the *emaj\_verify\_all()* function (§4.8.1), the administrator can perform the same checks on demand on all tables groups.

### 5.3.2 Event triggers

Installing E-Maj adds 2 event triggers of type “*sql\_drop*”:

- *emaj\_sql\_drop\_trg* blocks the drop attempts of:
  - ✓ any E-Maj object (log schema, log table, log sequence, log function and log trigger ),
  - ✓ any application table or sequence belonging to a table group in “*LOGGING*” state,
  - ✓ any schema containing at least one table or sequence belonging to a table group in “*LOGGING*” state.
- *emaj\_protection\_trg* blocks the drop attempts of the *emaj* extension itself and the main *emaj* schema.

Installing E-Maj also adds 1 event trigger of type “*table\_rewrite*”:

- *emaj\_table\_rewrite\_trg* blocks any structure change of application or log table.

It is possible to deactivate and reactivate these event triggers thanks to 2 functions: *emaj\_disable\_protection\_by\_event\_triggers()* and *emaj\_enable\_protection\_by\_event\_triggers()* (see §4.8.6).

However, the protections do not cover all risks. In particular, they do not prevent any tables or sequences renaming or any schema change. And some other DDL statements altering tables structure will not fire any trigger.

## 5.4 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*.

The *emaj\_hist* table structure is the following:

<b>Column</b>	<b>Type</b>	<b>Description</b>
<i>hist_id</i>	BIGSERIAL	serial number identifying a row in this history table

<i>hist_datetime</i>	TIMESTAMPTZ	recording date and time of the row
<i>hist_function</i>	TEXT	function associated to the traced event
<i>hist_event</i>	TEXT	kind of event
<i>hist_object</i>	TEXT	object related to the event (group, table or sequence)
<i>hist_wording</i>	TEXT	additional comments
<i>hist_user</i>	TEXT	role whose action has generated the event
<i>hist_txid</i>	BIGINT	identifier of the transaction that has generated the event

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

- *ADJUST\_GROUP\_PROPERTIES* adjust the *group\_has\_waiting\_changes* column content of the *emaj\_group* table
- *ALTER\_GROUP* tables group change
- *ALTER\_GROUPS* tables groups change
- *CLEANUP\_RLBK\_STATE* cleanup the state of recently completed rollback operations
- *COMMENT\_GROUP* comment set on a group
- *COMMENT\_MARK\_GROUP* comment set on a mark for a tables group
- *CONSOLIDATE\_RLBK\_GROUP* consolidate a logged rollback operation
- *CREATE\_GROUP* tables group creation
- *DBLINK\_OPEN\_CNX* open a *dblink* connection for a rollback operation
- *DBLINK\_CLOSE\_CNX* close a *dblink* connection for a rollback operation
- *DELETE\_MARK\_GROUP* mark deletion for a tables group
- *DISABLE\_EVENT\_TRIGGERS* deactivate event triggers
- *DROP\_GROUP* tables group suppression
- *EMAJ\_INSTALL* E-Maj installation or version update
- *ENABLE\_EVENT\_TRIGGERS* activate event triggers
- *FORCE\_DROP\_GROUP* tables group forced suppression
- *FORCE\_STOP\_GROUP* tables group forced stop
- *GEN\_SQL\_GROUP* generation of a psql script to replay updates for a tables group
- *GEN\_SQL\_GROUPS* generation of a psql script to replay updates for several tables groups
- *LOCK\_GROUP* lock set on tables of a group
- *LOCK\_GROUPS* lock set on tables of several groups
- *LOCK\_SESSION* lock set on tables for a rollback session
- *PROTECT\_GROUP* set a protection against rollbacks on a group
- *PROTECT\_MARK\_GROUP* set a protection against rollbacks on a mark for a group
- *PURGE\_HISTORY* delete from the *emaj\_hist* table the events prior the retention delay
- *RENAME\_MARK\_GROUP* mark rename for a tables group
- *RESET\_GROUP* log tables content reset for a group
- *ROLLBACK\_GROUP* rollback updates for a tables group
- *ROLLBACK\_GROUPS* rollback updates for several tables groups

- *ROLLBACK\_TABLE* rollback updates for one table
- *ROLLBACK\_SEQUENCE* rollback one sequence
- *SET\_MARK\_GROUP* mark set on a tables group
- *SET\_MARK\_GROUPS* mark set on several tables groups
- *SNAP\_GROUP* snap all tables and sequences for a group
- *SNAP\_LOG\_GROUP* snap all log tables for a group
- *START\_GROUP* tables group start
- *START\_GROUPS* tables groups start
- *STOP\_GROUP* tables group stop
- *STOP\_GROUPS* tables groups stop
- *UNPROTECT\_GROUP* remove a protection against rollbacks on a group
- *UNPROTECT\_MARK\_GROUP* remove a protection against rollbacks on a mark for a group

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

- *BEGIN*
- *END*
- *EVENT TRIGGERS DISABLED*
- *EVENT TRIGGERS ENABLED*
- *LOG DATA TABLESPACE CHANGED* tablespace for the log table modified
- *LOG INDEX TABLESPACE CHANGED* tablespace for the log index modified
- *LOG\_SCHEMA CREATED*
- *LOG\_SCHEMA DROPPED*
- *MARK DELETED*
- *NAMES PREFIX CHANGED* E-Maj names prefix modified
- *NOTICE* notice message issued by a rollback
- *SEQUENCE ADDED* *sequence added to a logging tables group*
- *SEQUENCE MOVED* *sequence moved from one group to another*
- *SEQUENCE REMOVED* *sequence removed from a logging tables group*
- *TABLE ADDED* *table added to a logging tables group*
- *TABLE MOVED* *table moved from one group to another*
- *TABLE REMOVED* *table removed from a logging tables group*
- *WARNING* warning message issued by a rollback

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

Two other internal tables keep traces of groups alter or rollback operations:

- *emaj\_alter\_plan* lists the elementary steps performed during the execution of *emaj\_alter\_group()* and related functions (Cf §4.3),
- *emaj\_rlbk\_plan* lists the elementary steps performed during the execution of *emaj\_rollback\_group()* and related functions (Cf §4.2.6 and 4.2.7).

When a tables group is started, using the *emaj\_start\_group()* function, or when old marks are deleted, using the *emaj\_delete\_before\_mark\_group()* function, the oldest events are deleted from *emaj\_hist* tables. The events kept are those not older than a parametrised

retention delay and not older than the oldest active mark and not older than the oldest uncompleted rollback operation. By default, the retention delay for events equals 1 year. But this value can be modified at any time by inserting the *history\_retention* parameter into the *emaj\_param* table with a SQL statement (see §5.1). The same retention applies to the tables that log elementary steps of tables groups alter or rollback operations.

## 5.5 IMPACTS ON INSTANCE AND DATABASE ADMINISTRATION

---

### 5.5.1 Stopping and restarting the instance

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

#### 5.5.1.1 General rule

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

If a transaction with table updates were not committed at instance stop, it would be rolled back during the recovery phase of the instance 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,...

#### 5.5.1.2 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 why application sequences are processed at the very end of the rollback operations (See §4.2.6). (For the same reason, at set mark time, application sequences are processed at the beginning of the operation.)

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

### 5.5.2 Saving and restoring the database



Using E-Maj allows a reduction in 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 databases on an external support.

### 5.5.2.1 File level saves and restores

When saving or restoring instances at file level, it is essential to save or restore ALL instance files, including those stored on dedicated tablespaces.

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.

### 5.5.2.2 Logical saves and restores of entire database

To properly save and restore a database with E-Maj, using *pg\_dump*, and *psql* or *pg\_restore*, it is essential that both source and restored databases use the **same E-Maj version**. If this is not the case, the content of some technical tables may be not synchronised with their structure. Reading the row of key '*emaj\_version*' in the *emaj.emaj\_param* table may help in knowing the version of an E-Maj extension created in a database.

Regarding stopped tables groups (in *IDLE* state), as log triggers are disabled and the content of related log tables is meaningless, there is no action required 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.

Restoring the database structure generates 2 error messages reporting that the *\_emaj\_protection\_event\_trigger\_fnct()* function and the *emaj\_protection\_trg* event trigger already exist:

```
...
ERROR:  function "_emaj_protection_event_trigger_fnct" already exists with same
argument types
...
ERROR:  event trigger "emaj_protection_trg" already exists
...
```

This message display is normal and does not indicate a defective restore. Indeed, both objects are created with the extension and are then detached from it, so that the trigger can block any attempt of the extension drop. As a result, the *pg\_dump* tool saves them as independent objects. And when restoring, these objects are created twice, first with the *emaj* extension creation, and then as independent objects, this second attempt generating both error messages.

### 5.5.2.3 Logical save and restore of partial database

With *pg\_dump* and *pg\_restore* tools, database administrators can perform on 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.

## 5.5.3 Data load

Beside using *pg\_restore* or *psql* with files produced by *pg\_dump*, it is possible to efficiently load large amounts of data with the *COPY* SQL verb or the *|copy psql* meta-command. In both cases, this data loading fires *INSERT* triggers, among them the E-Maj log trigger. Therefore, there is no constraint to use *COPY* or *|copy* in E-Maj environment.

With other loading tools, it is important to check that triggers are effectively fired for each row insertion.

## 5.5.4 Tables reorganisation

### 5.5.4.1 Reorganisation of application tables

Application tables protected by E-Maj can be reorganised using the SQL *CLUSTER* command. Whether or not log triggers are enabled, the organisation process has no impact on log tables content.

### 5.5.4.2 Reorganisation of E-Maj tables

The index corresponding to the primary key of each table from E-Maj schemas (neither log tables nor technical tables) is declared “cluster”.



So using E-Maj may have an operational impact regarding the execution of *CLUSTER* SQL commands at database level.

When E-Maj is used in continuous mode (with deletion of oldest marks instead of regular tables groups stop and restart), it is recommended to regularly reorganize E-Maj log tables. This reclaims unused disk space following mark deletions.

### 5.5.5 Using E-Maj with replication

#### 5.5.5.1 Integrated physical replication

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


However, because of the way PostgreSQL manages sequences, the sequences' current values may be a little forward on slave instances than on the master instance. For E-Maj, this may lightly overestimate the number of log rows in general statistics. But there is no consequence on the data integrity.

#### 5.5.5.2 Integrated logical replication

Starting with version 10, PostgreSQL includes logical replication mechanisms. The replication granularity is the table. The *publication* object used in logical replication is quite close to the E-Maj tables group concept, except that a *publication* cannot contain sequences.


Several cases have to be examined.

#### Replication of application tables managed by E-Maj




An application table that belongs to a tables group can be replicated. The effect of any rollback operation that may occur would be simply replicated on subscriber side, as long as no filter has been applied on replicated SQL verbs types.

## Replication of application tables with E-Maj activated on subscriber side



It is possible to include an application table into a tables group, with updates coming from a logical replication flow. But all E-Maj operations (starting/stopping the group, setting marks,...) must of course be executed on the subscriber side. An E-Maj rollback operation can be launched once the replication flow has been stopped (to avoid updates conflicts). But then, tables on both publisher and subscriber sides are not coherent anymore.


## Replication of E-Maj log tables



It is technically possible to replicate an E-Maj log table (once found a way to get the DDL that creates the log table – using *pg\_dump* for instance). This allows to duplicate or concentrate logs content on another server. But the replicated log table can only be used for log **auditing**. As log sequences or *TRUNCATE* verbs are not replicated, these logs cannot be used for other purposes.



## Replication of application tables and E-Maj log tables



Application tables and log tables can be simultaneously replicated. But as seen previously, these replicated logs can only be used for **auditing** purpose. E-Maj rollback operations can only be executed on publisher side.

### 5.5.5.3 Other replication solutions

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

## 5.6 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 the PostgreSQL instance.

The date and time of each update or each mark is recorded. But nothing other than 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.

## 5.7 PERFORMANCE

---

### 5.7.1 Updates recording overhead

Recording updates in E-Maj log tables has necessarily an impact on the duration of these updates. The global impact of this log on a given processing depends on numerous factors. Among them:

- the part that the update activity represents on the global processing,
- the intrinsic performance characteristics of the storage subsystem that supports log tables.

However, the E-Maj updates recording overhead is generally limited to a few per-cents. But this overhead must be compared to the duration of potential intermediate saves avoided with E-Maj.

### 5.7.2 E-Maj rollback duration

The duration of an E-Maj rollback depends on several factors, like:

- the number of updates to cancel,
- the intrinsic characteristics of the server and its storage material and the load generated by other activities hosted on the server,
- triggers or foreign keys on tables processed by the rollback operation,
- contentions on tables at lock set time.

To get an order of magnitude of an E-Maj rollback duration, it is possible to use the *emaj\_estimate\_rollback\_group()* and *emaj\_estimate\_rollback\_groups()* functions (See §4.6.3).

### 5.7.3 Optimizing E-Maj operations

Here are some advice to optimize E-Maj operations:

#### 5.7.3.1 Use tablespaces

Creating tables into tablespaces located in dedicated disks or file systems is a way to more efficiently spread the access to these tables. To minimize the disturbance of application tables access by log tables access, the E-Maj administrator has two ways to use tablespaces for log tables and indexes location.

By setting a specific default tablespace for the session before the tables groups creation, log tables and indexes are created by default into this tablespace, without any additional action. (See §Erreur : source de la référence non trouvée)

But through parameters set into the *emaj\_group\_def* table, it is also possible to specify a tablespace to use for any log table or log index. (See §4.2.2.3)

### 5.7.3.2 Declare foreign keys as DEFERRABLE

Foreign keys can be explicitly declared as *DEFERRABLE* at creation time. If a foreign key is declared *DEFERRABLE* and no *ON DELETE* or *ON UPDATE* clause is used, this foreign key is not dropped at the beginning and recreated at the end of an E-Maj rollback operation. The foreign key checks of updated rows are just deferred to the end of the rollback function execution, once all log tables are processed. This generally greatly speeds up the rollback operation.

### 5.7.3.3 Modify memory parameters

Increasing the value of the *work\_mem* parameter when performing an E-Maj rollback may bring some performance gains.

If foreign keys have to be recreated by an E-Maj rollback operation, increasing the value of the *maintenance\_work\_mem* parameter may also help.

If the E-Maj rollback functions are directly called in SQL, these parameters can be previously set at session level, with statements like:

```
SET work_mem = <value>;  
SET maintenance_work_mem = <value>;
```

If the E-Maj rollback operations are executed by a web client, it is also possible to set these parameters at function level, as *superuser*.

```
ALTER FUNCTION emaj._rlbk_tbl(emaj.emaj_relation, BIGINT, BIGINT, INT,  
BOOLEAN) SET work_mem = <value>;  
ALTER FUNCTION emaj._rlbk_session_exec(INT, INT) SET maintenance_work_mem  
= <value>;
```

## 5.8 USAGE LIMITS

---

The E-Maj extension usage has some limits.

- The minimum required PostgreSQL version is 9.5..
- All tables belonging to a “*rollbackable*” tables group must have an explicit *PRIMARY KEY*.
- *UNLOGGED* and *WITH OIDS* tables can only be members of “*audit\_only*” tables groups.
- *TEMPORARY* tables are not supported by E-Maj.
- 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. A trigger, created by E-Maj, blocks any *TRUNCATE* statement on any table belonging to a tables group in logging state.
- Using a global sequence for a database leads to a limit in the number of updates that E-Maj can manage throughout its life. This limit equals  $2^{63}$ , about  $10^{19}$  (but only  $10^{10}$  on oldest platforms), which still allow to record 10 million updates per second (100 times the best performance benchmarks results in 2012) during ... 30,000 years (or at worst 100 updates per second during 5 years). Would it be necessary to reset the global sequence, the E-Maj extension would just have to be un-installed and re-installed.
- 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 statement execution on the way E-Maj works, depending on the kind of executed operation.

- If a new table were created, it would be unable to enter into a group's definition until this group is 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 its 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 necessarily generate any error during further log recording. But the checks that E-Maj performs would block any attempt to set a new mark or 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 set a new mark or 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.

## 5.9 USER'S RESPONSIBILITY

---

### 5.9.1 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 operation are really included in a single tables group.

### 5.9.2 Appropriate call of main functions

*emaj\_start\_group()*, *emaj\_set\_mark\_group()*, *emaj\_rollback\_group()* and *emaj\_logged\_rollback\_group()* functions (and their related multi-groups 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 a stable point in the life of these tables. He must also take care of warning messages that may be reported by E-Maj rollback functions.

### 5.9.3 Management of application triggers

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

If the trigger simply adjusts the content of the row to insert or update, the logged data will contain the final columns value. So in case of rollback, the log table contains the right data to process. And as the trigger is by default automatically disabled at rollback time, the trigger cannot disturb the rollback processing

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

- if the updated table belongs to the same tables group, the automatic trigger disabling and the rollback of both tables will let them in the expected state,
- 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. If needed, the *emaj\_ignore\_app\_trigger()* function can be used to not disable the trigger at rollback time (see §4.8.3). But merely deactivating the trigger may not be sufficient and some other actions may be required.

For more complex triggers, it is essential to perfectly understand their impacts on E-Maj rollbacks and take any appropriate measure at rollback time.

#### 5.9.4 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 EMAJ\_WEB

---

A web application, *Emaj\_web*, makes E-Maj use much easier.

For the records, a plugin for *phpPgAdmin* also existed. But it is not maintained any more since E-Maj 3.0.

### 6.1 OVERVIEW

---

*Emaj\_web* has borrowed to *phpPgAdmin* its infrastructure (browser, icon trails, database connection, management,...) and some useful functions like browsing the tables content or editing SQL queries.

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, all E-Maj objects are accessible.

It is then possible to:

- define or modify groups content,
- see the list of tables groups and perform any possible action, depending on groups state (create, drop, start, stop, set or remove a mark, rollback, add or modify a comment),
- 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 and see their content,
- monitor in progress rollback operations.

### 6.2 EMAJ\_WEB CLIENT INSTALLATION

---

#### 6.2.1 Prerequisite

*Emaj\_web* requires a web server with a php interpreter.

#### 6.2.2 Download

The *Emaj\_web* application can be downloaded from the following git repository:

[https://github.com/dalibo/emaj\\_web](https://github.com/dalibo/emaj_web)

#### 6.2.3 Application configuration



The configuration is centralized into a single file: *emaj\_web/conf/config.inc.php*. It contains the general parameters of the applications, and the description of the PostgreSQL instances connections.

When the number of instances is large, it is possible to split them into “instances groups”. A group can contain instances or other instance groups.

In order to submit batch rollbacks (i.e. without blocking the use of the browser while the rollback operation is in progress), it is necessary to specify a value for two configuration parameters:

- `$conf['psql_path']` defines the access path of the *psql* executable file,
- `$conf['temp_dir']` defines a temporary directory that rollback functions can use.

The distributed *emaj\_web/conf/config.inc.php-dist* file can be used as a configuration template.

## 6.3 USING EMAJ\_WEB

---

### 6.3.1 Access to Emaj\_web and databases

Accessing Emaj\_web in a browser displays the welcome page.

To sign in to a database, select the target instance in the left browser or in the ‘*servers*’ tab, and fill the connection identifier and password. Several connections can remain opened simultaneously.

Once connected to a database where the emaj extension has been installed, the user interacts with the extension, depending on the role it owns (super-user, *emaj\_adm* or *emaj\_viewer*).

On the left, the browser tree shows all the configured instances, that may be split into instances groups, and all the databases they contain. By unfolding a database object, the user reaches the E-Maj tables groups and the existing schemas.

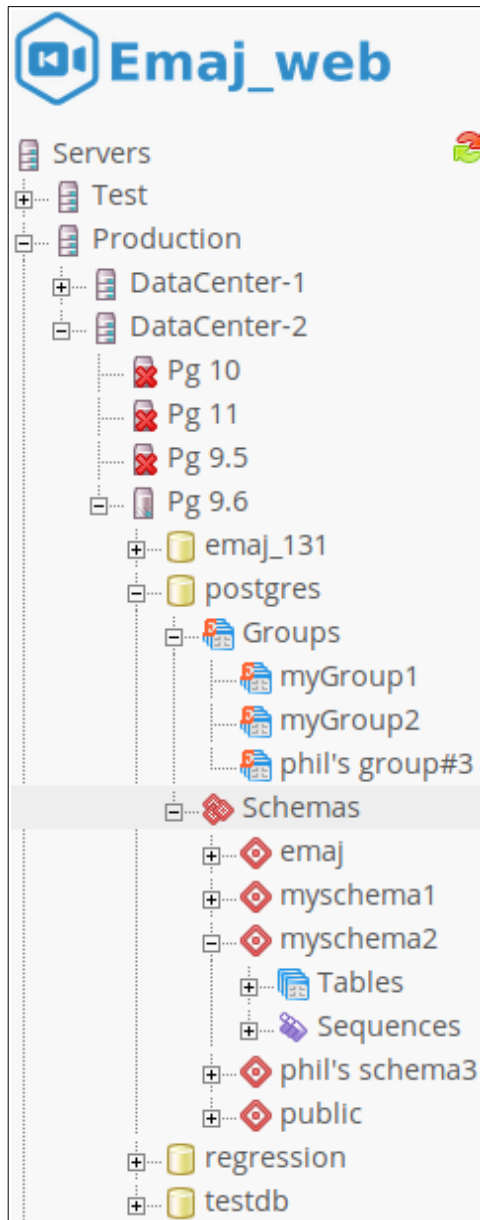


Figure 1 – The browser tree.

### 6.3.2 Tables groups list

By selecting a database, the user reaches a page that lists all tables groups created in this database.

Connection: localhost:5496 - role "postgres" SQL | History | Logout English

Emaj\_web > Pg 9.6 > postgres

Groups Groups configuration Rollback operations Schemas E-Maj

### Tables groups in "LOGGING" state

	Group	Creation date/time	# tables	# sequences	Type	# marks	Actions						Comment
<input type="checkbox"/>	myGroup1	21/02/2019 18:45:01	5	1		3	Set a mark	Protect	Rollback	Stop	Set a comment		Useless comment!
<input type="checkbox"/>	myGroup2	21/02/2019 18:45:01	4	2		4	Set a mark	Protect	Rollback	Stop	Set a comment		

Select Actions on selected objects (0)

All / Visible / None / Invert

### Tables groups in "IDLE" state

	Group	Creation date/time	# tables	# sequences	Type	# marks	Actions						Comment
<input type="checkbox"/>	phil's group#3	21/02/2019 18:45:02	2	1		0	Start	Reset	Set a comment				Drop

Select Actions on selected objects (0)

All / Visible / None / Invert

### Configured but not yet created tables groups

No tables group is currently configured but not created.

To create a (another) tables group, go first to the groups configuration tab.  
Alternatively, you can create an empty group, then add tables and sequences into it, and apply the configuration change.

Figure 2 – Tables groups list.

This page displays three lists:

- the tables groups in *LOGGING* state
- the tables groups in *IDLE* state,
- the tables groups referenced into the *emaj\_group\_def* table but not yet created.

For each created tables group, the following attributes are displayed:

- its creation date and time,
- the number of application tables and sequences it contains,
- its type ("*ROLLBACKABLE*" or "*AUDIT\_ONLY*", protected against rollback or not),
- the number of marks it owns,
- its associated comment, if any.

For each tables group, several buttons are available so that the user can perform any possible action, depending on the group state.

### 6.3.3 Some details about the user interface

The page headers contain:

- some information regarding the current connection,
- 3 links to reach the SQL statements editor, the history of submitted statements and to logout the current connection,
- a combo box to select the language used by the user interface,
- a breadcrumb trail,
- and a button allowing to directly go to the page bottom.

The user can navigate in Emaj\_web functions using two icon bars: one for the general purpose functions and the other for the functions concerning a single tables group.




Figure 3 – Main icons bar.




Figure 4 – Tables groups icons bar.

For *emaj\_viewer* roles, some icons are not visible.

On some tables, it is possible to dynamically sort displayed rows, using small vertical arrows on the right of column titles.

On some tables too, an icon located at the left of the header row, let show or hide input fields that can be used to dynamically filter displayed rows.

Tables groups in "LOGGING" state ⓘ														
⌵	Group ⌵	Creation date/time ⌵	# tables ⌵	# sequences ⌵	Type ⌵	# marks ⌵	Actions					Comment ⌵		
Reset	my					>2								
<input type="checkbox"/>	myGroup1	21/02/2019 18:45:01	5	1		3	Set a mark	Protect		Rollback	Stop	Set a comment		Useless comment!
<input type="checkbox"/>	myGroup2	21/02/2019 18:45:01	4	2		4	Set a mark	Protect		Rollback	Stop	Set a comment		

Figure 5 – Filtering the tables groups in logging state.

Here, only tables groups whose name contains “my” and having more than 2 marks are displayed, sorted in descending order by number of tables.

Some tables allow to perform actions on several objects at once. In this case, the user selects the objects with the checkboxes on the first column of the table and choose the action to perform among the available buttons under the table.

### 6.3.4 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. This page is also accessible with the “*Properties*” icon of the groups bar and through the left browsing tree.

Connection: localhost:5496 - role "postgres" SQL | History | Logout English

Emaj\_web > Pg 9.6 > postgres > myGroup1

Properties Log statistics Content

#### Tables group "myGroup1" properties

State	Creation date/time	Type	# tables	# sequences	# marks	Log size
	2019-02-21 18:45:01.885922+01		5	1	3	144 kB

Comment: *Useless comment!*

Set a mark Protect Stop Set a comment

#### Tables group "myGroup1" marks

	Mark	State	Date/Time	# row changes	Cumulative changes	Actions						Comment	
<input type="checkbox"/>	MARK3		2019-02-21 18:45:02.211977+01	0	0	Rollback	Rename	Delete	First mark	Protect		Set a comment	
<input type="checkbox"/>	MARK2		2019-02-21 18:45:02.181251+01	7	7	Rollback	Rename	Delete	First mark	Protect		Set a comment	End of 1st program
<input type="checkbox"/>	MARK1		2019-02-21 18:45:02.102865+01	19	26	Rollback	Rename	Delete		Protect		Set a comment	

Select Actions on selected objects (0)

All / Visible / None / Invert Delete

Figure 6 – Details of a tables group

A first line repeats information already displayed on the groups list (number of tables and sequences, type, state and number of marks). It also shows the disk space used by its log tables.

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

Next is a set of buttons to execute actions depending on the group's state.

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

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

- its state (active or not, protected against rollback or not),
- 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 was set,
- the comment associated to the mark, if it exists.

For each mark, several buttons are available to perform the actions permitted by the mark's state.

### 6.3.5 Statistics

Using the “*Log statistics*” tab of the group’s bar, one gets statistics about updates recorded into the log tables for the selected tables group.

Two types of statistics can be produced:

- some estimates about the number of updates per table, recorded between two marks or between one mark and the current situation,
- a precise numbering of updates per tables, per statement type (*INSERT/UPDATE/DELETE/TRUNCATE*) and role.

If the end of the range corresponds to the current situation, a check box allows one to request a rollback simulation to the selected mark in order to quickly get an approximate duration of this operation.

The figure below shows an example of detailed statistics.

Connection: localhost:5496 - role "postgres" SQL | History | Logout | English

Emaj\_web > Pg 9.6 > postgres > myGroup1

Properties
Log statistics
Content

**Statistics from E-Maj logs for the group "myGroup1"**

From MARK2 (2019-02-21 18:45:02.181251+01) To MARK3 (2019-02-21 18:45:02.211977+01)

Simulate Rollback  Estimates Detailed stats ⚠

**Table changes between the mark "MARK2" and the mark "MARK3" for the tables group "myGroup1"**

# tables	# row changes	# INSERT	# UPDATE	# DELETE	# TRUNCATE	# roles	Roles
4	7	5	1	1	0	1	postgres

	Schema	Table	First mark	First mark date/time	Last mark	Last mark date/time	Role	SQL verb	# row changes	Actions
	myschema1	mytbl1	MARK2	2019-02-21 18:45:02.197825+01	MARK3	2019-02-21 18:45:02.227135+01	postgres	DELETE	1	<a href="#">SQL</a>
	myschema1	mytbl1	MARK2	2019-02-21 18:45:02.197825+01	MARK3	2019-02-21 18:45:02.227135+01	postgres	INSERT	2	<a href="#">SQL</a>
	myschema1	mytbl2	MARK2	2019-02-21 18:45:02.197825+01	MARK3	2019-02-21 18:45:02.227135+01	postgres	INSERT	1	<a href="#">SQL</a>
	myschema1	mytbl2b	MARK2	2019-02-21 18:45:02.197825+01	MARK3	2019-02-21 18:45:02.227135+01	postgres	INSERT	1	<a href="#">SQL</a>
	myschema1	myTbl3	MARK2	2019-02-21 18:45:02.197825+01	MARK3	2019-02-21 18:45:02.227135+01	postgres	INSERT	1	<a href="#">SQL</a>
	myschema1	myTbl3	MARK2	2019-02-21 18:45:02.197825+01	MARK3	2019-02-21 18:45:02.227135+01	postgres	UPDATE	1	<a href="#">SQL</a>

*Figure 7 – Detailed statistics about updates recorded between two marks*

The displayed page contains a first line returning global counters.

On each line of the statistics table, the user can click on a “SQL” button to easily look at the log tables content. A click on this button opens the SQL editor window and proposes the statement displaying the content of the log table that corresponds to the selection (table, time frame, role, statement type). The user can modify this suggested statement before executing it to better fit his needs.

Connection: localhost:5496 - role "postgres" SQL | History | Logout | English

Emaj\_web > Pg 9.6 > postgres > myGroup1

Properties Log statistics Content

**Statistics from E-Maj logs for the group "myGroup1"**

From:  To:

Simulate Rollback

**Table changes between the mark "MARK1" and the current state for the tables group "myGroup1"**

Estimates	
# tables	# row changes
5	26

Rolling the tables group "myGroup1" back to the mark "MARK1" would take about 00:00:01.

	Schema	Table	First mark	First mark date/time	Last mark	Last mark date/time	# row changes	Actions
	myschema1	mytbl1	MARK1	2019-02-21 18:45:02.158836+01			7	<a href="#">SQL</a>
	myschema1	mytbl2	MARK1	2019-02-21 18:45:02.158836+01			3	<a href="#">SQL</a>
	myschema1	mytbl2b	MARK1	2019-02-21 18:45:02.158836+01			3	<a href="#">SQL</a>
	myschema1	myTbl3	MARK1	2019-02-21 18:45:02.158836+01			12	<a href="#">SQL</a>
	myschema1	mytbl4	MARK1	2019-02-21 18:45:02.158836+01			1	<a href="#">SQL</a>

Figure 8 – Result of the rollback simulation, with the estimated number of updates for each table.

The displayed page contains a first part indicating the number of tables and sequences concerned by the rollback operation to this mark, and an estimate of the operation duration.

### 6.3.6 Tables group content

Using the “Content” tab of the group’s bar, it is possible to get a summary of a tables group content.

For each table and sequence belonging to the group, the displayed table shows the characteristics configured into the *emaj\_group\_def* table, as well as the disk space used by the log table and its index.



Connection: localhost:5496 - role "postgres" SQL | History | Logout | English

Emaj\_web > Pg 9.6 > postgres > myGroup1

Properties
Log statistics
Content

**Content of the tables group "myGroup1"**

	Type	Schema	Name	Priority	Log schema	Log tablespace	Log index tablespace	Objects name prefix	Log size	Log size
		myschema1	mytbl1	20	emaj			myschema1_mytbl1	32768	32 kB
		myschema1	mytbl2		emaj			myschema1_mytbl2	32768	32 kB
		myschema1	mytbl2b		emaj			myschema1_mytbl2b	24576	24 kB
		myschema1	myTbl3	10	emaj			myschema1_myTbl3	24576	24 kB
		myschema1	myTbl3_col31_seq	1						
		myschema1	mytbl4	20	emaj			myschema1_mytbl4	32768	32 kB

*Figure 9 – A tables group’s content.*

### 6.3.7 Tables groups configuration

By selecting the “*Groups configuration*” tab of the main bar, the user reaches the function that manages the tables groups content configuration.

The upper part of the page lists the existing schemas (except schemas dedicated to E-Maj). By selecting a schema, the list of its tables and sequences appears.

Connection: localhost:5496 - role "postgres" SQL | History | Logout English

Emaj\_web > Pg 9.6 > postgres

Groups Groups configuration Rollback operations Schemas E-Maj

### Application schemas

Y	Schema	Owner	Comment
	myschema1	postgres	
	myschema2	postgres	
	phil's schema3	postgres	
	public	postgres	standard public schema

### Tables and sequences in schema "myschema1"

Y	Type	Schema	Name	Actions	Group	Priority	Log schema suffix	Objects name prefix	Log tablespace	Log Index tablespace
<input type="checkbox"/>	Table	myschema1	myTbl3	Update Remove	myGroup1	10				
<input type="checkbox"/>	Sequence	myschema1	myTbl3_col31_seq	Update Remove	myGroup1	1				
<input type="checkbox"/>	Table	myschema1	mytbl1	Update Remove	myGroup1	20				
<input type="checkbox"/>	Table	myschema1	mytbl2	Update Remove	myGroup1					
<input type="checkbox"/>	Table	myschema1	mytbl2b	Update Remove	myGroup1					
<input type="checkbox"/>	Sequence	myschema1	mytbl2b_col20_seq	Assign						
<input type="checkbox"/>	Table	myschema1	mytbl4	Update Remove	myGroup1	20				

Select Actions on selected objects (0)  
 All / Visible / None / Invert / Not assigned Assign Remove

Figure 10 – Tables groups content configuration.

The user can then view or modify the tables and sequences assignment into tables groups. All these actions update the *emaj\_group\_def* table used for the tables groups creation.

The following information is listed for each table or sequence:

- its type
- the tables group it belongs to, if any,
- the following attributes of the table or sequence in the *emaj\_group\_def* table, if assigned: (see §4.2.2):
  - the priority level in the group,
  - the suffix that defines log schema,
  - the prefix used to build the E-Maj object names for this table,
  - the optional tablespace name for the log table,
  - the optional tablespace name for the log table's index,
- its owner,
- the tablespace it belongs to, if any,
- the associated comment in the database.

The schemas list and the tables and sequences list also display the objects that are known in the *emaj\_group\_def* table but don't exist in the database. These objects are identified with a “!” icon in the first column of each list.

With buttons, it is possible to:

- assign a table or a sequence to a new or an already known tables group,
- modify the properties of a table or a sequence inside its tables group,
- remove a table or a sequence from its tables group.

Note that any change applied in the *emaj\_group\_def* table's content will only be effective when the concerned tables groups will be altered, or dropped and re-created.

### 6.3.8 Monitoring rollback operations

Using the “*Rollback operations*” tab of the main bar, users can monitor the rollback operations. Three different lists are displayed:

- in progress rollback operations, with the characteristics of the rollback operations and estimates of the percentage of the operation already done and of the remaining duration,
- the latest completed operations,
- logged rollback operations that are consolidable.

For completed operations, the user can filter on a more or less deep history.

For each consolidable rollback, a button allows to effectively consolidate the operation.

Connection: localhost:5496 - role "postgres" SQL | History | Logout English

Emaj\_web > Pg 9.6 > postgres

Groups Groups configuration Rollback operations Schemas E-Maj

### In progress E-Maj rollbacks

Rlbk Id.	Groups	State	Rollback start	Current duration	Estimated remaining	% completed	Target mark	Mark set at	Logged ?	Nb sessions	Nb tables to process	se to
4	myGroup2	LOCKING			00:00:00.044023	0	MARK1	2019-02-21 18:45:02.158836+01	Yes	1	2	

### Completed E-Maj rollbacks

Display the :  3 most recent  completed since less than 24 hours [Filter](#)

Rlbk Id.	Groups	State	Rollback start	Rollback end	Duration	Target mark	Mark set at	Logged ?	Nb sessions	Nb process table
3	myGroup2	COMMITTED	2019-02-21 18:45:02.378441+01	2019-02-21 18:45:02.414558+01	00:00:00	MARK3	2019-02-21 18:45:02.211977+01	Yes	1	
2	myGroup2	COMMITTED	2019-02-21 18:45:02.325197+01	2019-02-21 18:45:02.371898+01	00:00:00	tmp_mark	2019-02-21 18:45:02.23706+01	No	1	
1	myGroup2	COMMITTED	2019-02-21 18:45:02.257901+01	2019-02-21 18:45:02.322943+01	00:00:00	MARK1	2019-02-21 18:45:02.102865+01	Yes	1	

### Consolidable E-Maj logged rollbacks

Group	Target mark	Mark set at	# row changes	Nb Intermediate marks	End rollback mark	Mark set at	Actions
myGroup2	MARK3	2019-02-21 18:45:02.211977+01	6	1	RLBK_MARK3_18.45.02.395_DONE	2019-02-21 18:45:02.378441+01	<a href="#">Consolidate</a>

Figure 11– Rollback operations monitoring.

### 6.3.9 Schemas content

The “Schemas” tab allows to select schemas in order to:

- list tables or sequences it contains,
- see the structure or the content of a particular table,
- see the properties of a particular sequence.

### 6.3.10 E-Maj environment state

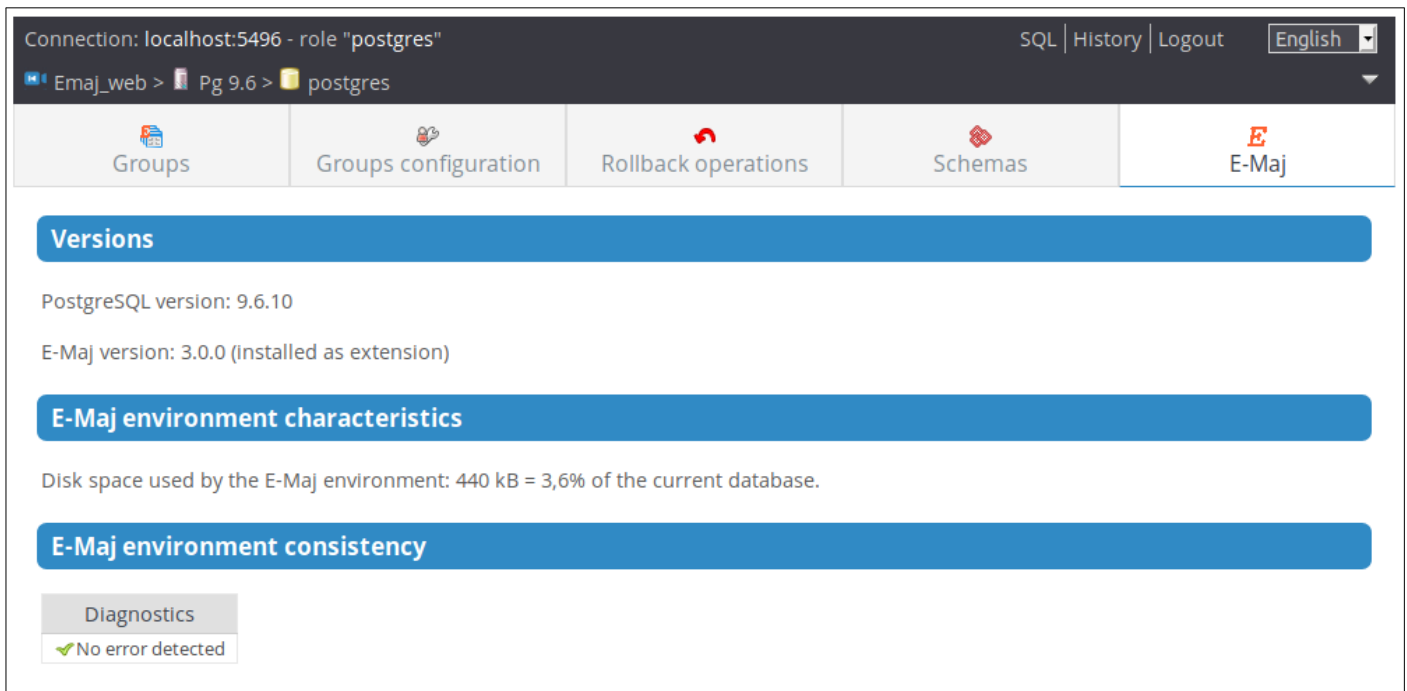
By selecting the “E-Maj” tab of the main bar, the user reaches an overview of the E-Maj environment state.

First, some items are displayed:

- the installed PostgreSQL and E-Maj versions,

- the disk space used by E-Maj (log tables, technical tables and their indexes), and the part of the global database space it represents.

Then, the environment integrity is checked; the result of the `emaj_verify_all()` function execution is displayed.



The screenshot shows the E-Maj web interface. At the top, it displays the connection information: "Connection: localhost:5496 - role 'postgres'". Below this, there are navigation tabs: "Groups", "Groups configuration", "Rollback operations", "Schemas", and "E-Maj". The "E-Maj" tab is selected. Underneath, there are three main sections:

- Versions**: Shows "PostgreSQL version: 9.6.10" and "E-Maj version: 3.0.0 (Installed as extension)".
- E-Maj environment characteristics**: Shows "Disk space used by the E-Maj environment: 440 kB = 3,6% of the current database."
- E-Maj environment consistency**: Shows a "Diagnostics" section with a green checkmark and the text "No error detected".

Figure 12 – The E-Maj environment state.

## 7 CONTRIBUTE TO THE E-MAJ DEVELOPMENT

---

Any contribution to the development and the improvement of the E-Maj extension is welcome. This chapter gives some information to make these contributions easier.

### 7.1 BUILD THE E-MAJ ENVIRONMENT

---

The E-Maj extension repository is hosted on the *github* site: <https://github.com/dalibo/emaj>

#### 7.1.1 Clone the E-Maj repository

So the first action to perform is to locally clone this repository on his/her own computer. This can be done by using the functionalities of the *github* web interface or by typing the shell command:

```
git clone https://github.com/dalibo/emaj.git
```

#### 7.1.2 Description of the E-Maj tree

So one has a full directory tree (except the web clients). It contains all directories and files described in the appendix 8.2, except the *doc* directory content that is separately maintained (see below).

The main directory also contains the following components:

- the *tar.index* file that is used to build the tarball of the E-Maj version distributed on *pgxn.org*
- the *docs* directory with all sources of the online documentation (see §7.4)
- in the *sql* directory:
  - the file *emaj--devel.sql*, source of the extension in its current version
  - the source of the previous version *emaj--<previous\_version>.sql*
  - a *emaj\_prepare\_emaj\_web\_test.sql* script that prepares an E-Maj environment to test the *Emaj\_web* client
- a *test* directory containing all components used to test the extension (see §7.3)
- a *tools* directory containing some ... tools.

#### 7.1.3 Setting tools parameters

The tools stored in the *tools* directory need some parameters to be set, depending on his/her own environment. A parameter system covers some tools. For the others, the *tools/README* file details the changes to apply.

### 7.1.3.1 Créating the `emaj_tools.env` file

The parameters that may be modified are grouped into the `tools/emaj_tools.env` file, which is called by `tools/emaj_tools.profile`.

The repository contains a file `tools/emaj_tools.env-dist` that may be used as a template to create the `emaj_tools.env` file.

The `emaj_tools.env` file must contain:

- the list of PostgreSQL versions that are supported by the current E-Maj version and for which a PostgreSQL instance exists for tests (`EMAJ_USER_PGVER` variable),
- for each PostgreSQL version used for the tests, 6 variables describing the location of binaries, the main directory of the related instance, the role and the ip-port to be used for the connection to the instance.

## 7.2 CODING

---

### 7.2.1 Versionning

The version currently under development is named *devel*.

Regularly and when it is justified, a new version is created. Its name has a X.Y.Z pattern.

The `tools/create_version.sh` shell script assists in creating this version. It is only used by the E-Maj maintainers. So its use is not described here.

### 7.2.2 Coding rules

Coding the `emaj--devel.sql` script must follow these rules:

- script structure: after some checks about the execution conditions that must be met, the objects are created in the following order: roles, enumerated types, sequences, tables (with their indexes and constraints), composite types, E-Maj parameters, low level functions, elementary functions that manage tables and sequences, functions that manage tables groups, general purpose functions, event triggers, grants, additional actions for the extensions. The script ends with some final operations.
- all objects are created in the `emaj` schema, except the `_emaj_protection_event_trigger_fnct()` function, created in the `public` schema,
- tables and sequences names are prefixed by `'emaj_'`
- functions names are prefixed by `'emaj_'` when they are usable by end users, or by `'_'` for internal functions,
- the internal tables and the functions callable by end users must have a comment,
- the language keywords are in upper case, objects names are in lower case,
- the code is indented with 2 space characters,

- lines must not contain tab characters, must not be longer than 140 characters long and must not end with spaces,
- in the functions structure, the code delimiters must contain the function name surrounded with a \$ character (or *\$do\$* for code blocks),
- variables names are prefixed with 'v\_' for simple variables or 'r\_' for *RECORD* type variables,
- the code must be compatible with all PostgreSQL versions supported by the current E-Maj version. When this is strictly necessary, the code may be differentiated depending on the PostgreSQL version.

A perl script, *tools/check\_code.pl* performs some checks on the code format of the script that creates the extension. It also detects unused variables. This script is directly called in non-regression tests scenarios (Cf §7.3).

### 7.2.3 Version upgrade script

E-Maj is installed into a database as an extension. The E-Maj administrator must be able to easily upgrade the extension version (cf §3.4). So an upgrade script is provided for each version, that upgrades from the previous version to the next version. It is named *emaj--<previous\_version>--devel.sql*.

The development of this script follows these rules:

- Develop/maintain the upgrade script at the same time as the main *emaj--devel.sql* script, so that the tests of a change include upgrade version cases,
- Apply the same coding rules as for the main script,
- As far as possible, ensure that the upgrade operation is able to process tables groups in logging state, without losing the capability to perform E-Maj rollbacks on marks set prior the version upgrade.

At the beginning of a version, the upgrade script is built using a template (the file *tools/emaj\_upgrade.template*).

As the development goes on, a perl script helps to synchronize the creation/deletion/replacement of functions. It compares the *emaj--devel.sql* script and the script that creates the previous version and updates the *emaj--<previous\_version>--devel.sql* script. To let it work properly, it is essential to keep both tags that frame the part of the script that describes functions.

After having adapted the parameters (see the *TOOLS/README* file), just submit:

```
perl tools/sync_fct_in_upgrade_script.pl
```

The other parts of the script must be coded manually. If the structure of an internal table is changed, the table content must be migrated (scripts for prior version upgrade can be used as examples).



## 7.3 TESTING

---

Through the rollback functions, the E-Maj extension updates database content. So the reliability is a key characteristics. For this reason, it is essential to pay a great attention to the tests.

### 7.3.1 Create PostgreSQL instances

The ideal is to be able to test E-Maj with all PostgreSQL versions that are supported by the extension (currently from version 9.5 to version 12).

The `tools/create_cluster.sh` script helps in creating a test instance. Its content may show the characteristics of the instance to create. It can also be executed (after parameters setting as indicated in `tools/README`):

```
tools/create_cluster.sh <PostgreSQL_major_version>
```

### 7.3.2 Install software dependancies

Testing the clients may require to install some additional software components:

- the *php* software, with its PostgreSQL interface,
- the *perl* software, with the *DBI* and *DBD::Pg* modules.

### 7.3.3 Execute non regression tests

A solid test environment is supplied in the repository. It contains:

- a test tool,
- test scenarios,
- expected results.

#### 7.3.3.1 The test scenarios

The test system contains 4 scenarios:

- a full standart scenario,
- the same scenario but installing the extension from the previous version with an immediate upgrade into the current version,
- the same scenario but installing the extension with the `emaj-devel.sql` script provided for cases when a “CREATE EXTENSION emaj” statement is not possible,
- a shorter scenario but with an upgrade into the current version while tables groups are in logging state.

These scenarios call *psql* scripts, all located into the *test/sql* directory. The scripts chain E-Maj function calls in different contexts, and SQL statements to prepare or check the results.

At the end of scripts, internal sequences are often reset, so that a single function call insertion does not produce impacts in the next scripts results.

The *psql* test scripts must be maintained in the same time as the extension source.

### 7.3.3.2 The expected results

For each *psql* script, the test tool produces a result file. These files are distinguished from a PostgreSQL version to another. They are located in the *test/<PostgreSQL\_version>/results* directory.

At the end of a run, the test tool compares these files with a reference located into the *test/<PostgreSQL\_version>/expected* directory.

Unlike for files in the *test/<PostgreSQL\_version>/results* directory, files in the *test/<PostgreSQL\_version>/expected* directory belong to the git repository. They must always remain consistent with the source of the extension and the *psql* test scripts.

### 7.3.3.3 The test tool

The test tool, *regress.sh*, combines all test functions.

Before using it, it is necessary to:

- have the PostgreSQL instances to be used already created and the *tools/emaj\_tools.env* file already setup
- manually create the *test/<PostgreSQL\_version>/results* directories.

The test tool can be launched with the command:

```
tools/regress.sh
```

As it starts with a copy of the *emaj.control* file into the *SHAREDIR/extension* directory of each configured PostgreSQL version, it may ask for the password of the Linux account to be able to execute *sudo* commands. It also automatically generates the *emaj-devel.sql* script used to create the extension with *psql*.

It then displays the list of test functions in a menu. Just enter the letter corresponding to the chosen test.

The test functions are:

- standart tests for each configured PostgreSQL version,

- the tests with the installation of the previous version followed by an upgrade,
- the tests with the installation of the version with the *emaj-devel.sql* script,
- the tests with an E-Maj version upgrade while tables groups are in logging state,
- tests chaining a database save with *pg\_dump* and a restore, with different PostgreSQL versions,
- a PostgreSQL upgrade version test using *pg\_upgrade* with a database containing the E-Maj extension.

It is important to execute the four first sets of tests for each E-Maj change.

### 7.3.3.4 Validate results

After having executed a *psql* script, *regress.sh* compares the outputs of the run with the expected outputs and reports the comparison result with the words 'ok' or 'FAILED'.

Here is an example of the display issued by the test tool (in this case with the scenario chaining the installation and a version upgrade, and with a detected difference):

```

Run regression test
===== dropping database "regression" =====
DROP DATABASE
===== creating database "regression" =====
CREATE DATABASE
ALTER DATABASE
===== running regression test queries =====
test install_upgrade      ... ok
test setup                ... ok
test create_drop         ... ok
test start_stop          ... ok
test mark                 ... ok
test rollback            ... ok
test misc                 ... ok
test alter                ... ok
test alter_logging       ... ok
test viewer              ... ok
test adm1                 ... ok
test adm2                 ... ok
test client               ... ok
test check                ... FAILED
test cleanup              ... ok

=====
 1 of 15 tests failed.
=====

The differences that caused some tests to fail can be viewed in the
file "/home/postgres/proj/emaj/test/11/regression.diffs". A copy of the test summary
that you see
above is saved in the file "/home/postgres/proj/emaj/test/11/regression.out".

```

When at least one script fails, it is important to closely analyze the differences, by reviewing the *test/<PostgreSQL\_version>/regression.diffs* file content, and check that the differences are directly linked to changes applied in the extension source code or in the test scripts.

Once the reported differences are considered as valid, the content of the *test/<PostgreSQL\_version>/result* directories must be copied into the *test/<PostgreSQL\_version>/expected* directories. A shell script processes all PostgreSQL versions in a single command:

```
sh tools/copy2Expected.sh
```

It may happen that some test outputs do not match the expected outputs, due to differences in the PostgreSQL behaviour from one run to another. Repeating the test allows to check these cases.

## 7.3.4 Test coverage

### 7.3.4.1 Functions test coverage

The PostgreSQL test instances are configured to count the functions executions. The *check.sql* test script displays the functions execution counters. It also displays E-Maj functions that have not been executed.

### 7.3.4.2 Error messages test coverage

A perl script extracts error and warning messages coded in the *sql/emaj--devel.sql* file. It then extracts the messages from the files of the *test/10/expected* directory. It finally displays error or warning messages that are not covered by tests.

The script can be run with the command:

```
perl tools/check_error_messages.pl
```

Some messages are known to not be covered by tests (for instance internal errors that are hard to reproduce). These messages, coded in the perl script, are excluded from the final report.

## 7.3.5 Evaluate the performances

The *tools/performance* directory contains some shell scripts helping in measuring performances. As the measurement results totally depend on the platform and the environment used, no reference results are supplied.

The scripts cover the following domains:

- *log\_overhead/pgbench.sh* evaluates the log mechanism overhead, using pgbench,
- *large\_group/large\_group.sh* evaluates the behaviour of groups containing a large number of tables,
- *rollback/rollback\_perf.sh* evaluates the E-Maj rollback performances with different tables profiles.

For all these files, some variables have to be configured at the beginning of the scripts.

## 7.4 DOCUMENTING

---

A *LibreOffice* format documentation is managed by the maintainers. It has its own *github* repository: *emaj\_doc*. Thus the *doc* directory of the main repository remains empty.

The online documentation is managed by *sphinx*. It is located in the *docs* directory.

To install *sphinx*, refer to the *docs/README.rst* file.

The documentation exists in two languages, English and French. Depending on the languages, document sources are located in */docs/en* and */docs/fr*. These documents are in *ReStructured Text* format.

To compile the documentation for a language, set the current directory to *docs/<language>* and execute the command:

```
make html
```

When there is no compilation error anymore, the documentation becomes available locally on a browser, by opening the *docs/<language>/\_build/html/index.html* file.

The documentation on the *readthedocs.org* site is automatically updated as soon as the main *github* repository is updated.

## 7.5 SUBMITTING A PATCH

---

Patches can be proposed to the E-Maj maintainers through *Pull Requests* on the *github* site.

Before submitting a patch, it may be useful to create an *issue* on *github*, in order to start a discussion with the maintainers and help in working on the patch.

## 8 APPENDIX

### 8.1 E-MAJ FUNCTIONS LIST

E-Maj functions that are available to users are listed in alphabetic order below. They are all callable by roles having *emaj\_adm* privileges. The chart also specifies those callable by *emaj\_viewer* roles.

Functions	Input parameters	Output data	Callable by emaj_viewer	Ref.
<b>emaj_alter_group</b>	group TEXT [ mark TEXT ]	# tables.and.seq INT		§ 4.3
<b>emaj_alter_groups</b>	groups.array TEXT [ mark TEXT ]	# tables.and.seq INT		§ 4.3
<b>emaj_cleanup_rollback_state</b>	-	# rollback INT		§ 14.8.5
<b>emaj_comment_group</b>	group TEXT comment TEXT	-		§ 4.4.2
<b>emaj_comment_mark_group</b>	group TEXT mark TEXT comment TEXT	-		§ 4.5.1
<b>emaj_consolidate_rollback_group</b>	group TEXT end.rollback.mark TEXT	# tables.and.seq INT		§4.4.6
<b>emaj_create_group</b>	group TEXT [ is.rollbackable BOOLEAN ] [ is.empty BOOLEAN ]	#.tables.and.seq INT		§ 4.2.3
<b>emaj_delete_before_mark_group</b>	group TEXT mark TEXT	#.deleted.marks INT		§ 4.5.5
<b>emaj_delete_mark_group</b>	group TEXT mark TEXT	1 INT		§ 4.5.4
<b>emaj_detailed_log_stat_group</b>	group TEXT start.mark TEXT end.markTEXT	SETOF emaj_detailed_log_stat_type	Yes	§ 4.6.2
<b>emaj_detailed_log_stat_groups</b>	groups.array TEXT[] start.mark TEXT end.mark TEXT	SETOF emaj_log_stat_type	Yes	§ 4.6.2
<b>emaj_disable_protection_by_event_triggers</b>	-	# triggers INT		§ 4.8.6
<b>emaj_drop_group</b>	group TEXT	#.tables.and.seq INT		§ 4.2.9
<b>emaj_enable_protection_by_event_triggers</b>	-	# triggers INT		§ 4.8.6
<b>emaj_estimate_rollback_group</b>	group TEXT mark TEXT	duration INTERVAL	Yes	§ 4.6.3

Functions	Input parameters	Output data	Callable by emaj_viewer	Ref.
<b>emaj_estimate_rollback_groups</b>	groups.array TEXT[] mark TEXT	duration INTERVAL	Yes	§ 4.6.3
<b>emaj_force_drop_group</b>	group TEXT	#.tables.and.seq INT		§4.4.5
<b>emaj_force_stop_group</b>	group TEXT	#.tables.and.seq INT		§ 4.4.4
<b>emaj_gen_sql_group</b>	group TEXT start.mark TEXT end.mark TEXT output.file.path TEXT [ tables.seq.array TEXT[] ]	#.gen.statements BIGINT		§ 4.7.3
<b>emaj_gen_sql_groups</b>	groups.array TEXT[] start.mark TEXT end.mark TEXT output.file.path TEXT [ tables.seq.array TEXT[] ]	#.gen.statements BIGINT		§ 4.7.3
<b>emaj_get_consolidable_rollbacks</b>	-	SETOF emaj_consolidable_rol lback_type	Yes	§ 4.4.7
<b>emaj_get_current_log_table</b>	schema TEXT table TEXT	(log.schema TEXT, log.table TEXT)	Yes	§ 4.8.2
<b>emaj_get_previous_mark_group</b>	group TEXT date.time TIMESTAMPTZ	mark TEXT	Yes	§ 4.5.2
<b>emaj_get_previous_mark_group</b>	group TEXT mark TEXT	mark TEXT	Yes	§ 4.5.2
<b>emaj_ignore_app_trigger</b>	action TEXT schema TEXT table TEXT trigger TEXT	#.triggers INT		§ 4.8.3
<b>emaj_log_stat_group</b>	group TEXT start.mark TEXT end.mark TEXT	SETOF emaj_log_stat_type	Yes	§ 4.6.1
<b>emaj_log_stat_groups</b>	groups.array TEXT[] start.mark TEXT end.mark TEXT	SETOF emaj_log_stat_type	Yes	§ 4.6.1
<b>emaj_logged_rollback_group</b>	group TEXT mark TEXT is_alter_group_allowed BOOLEAN	SETOF (severity TEXT, message TEXT)		§ 4.2.7
<b>emaj_logged_rollback_groups</b>	groups.array TEXT[] mark TEXT is_alter_group_allowed BOOLEAN	SETOF (severity TEXT, message TEXT)		§ 4.2.7
<b>emaj_protect_group</b>	group TEXT	0/1 INT		§ 4.4.3
<b>emaj_protect_mark_group</b>	group TEXT mark TEXT	0/1 INT		§ 4.5.6



Functions	Input parameters	Output data	Callable by emaj_viewer	Ref.
<b>emaj_rename_mark_group</b>	group TEXT mark TEXT new.name TEXT	-		§ 4.5.3
<b>emaj_reset_group</b>	group TEXT	#.tables.and.seq INT		§ 4.4.1
<b>emaj_rollback_activity</b>	-	SETOF emaj_rollback_activity _type	Yes	§ 4.8.4
<b>emaj_rollback_group</b>	group TEXT mark TEXT is_alter_group_allowed BOOLEAN	SETOF (severity TEXT, message TEXT)		§ 4.2.6
<b>emaj_rollback_groups</b>	groups.array TEXT[] mark TEXT is_alter_group_allowed BOOLEAN	SETOF (severity TEXT, message TEXT)		§ 4.2.6
<b>emaj_set_mark_group</b>	group TEXT [ mark TEXT ]	#.tables.and.seq INT		§ 4.2.5
<b>emaj_set_mark_groups</b>	groups.array TEXT[] [ mark TEXT ]	#.tables.and.seq INT		§ 4.2.5
<b>emaj_snap_group</b>	group TEXT directory TEXT copy.options TEXT	#.tables.and.seq INT		§ 4.7.1
<b>emaj_snap_log_group</b>	group TEXT start.mark TEXT end.mark TEXT directory TEXT copy.options TEXT	#.tables.and.seq INT		§ 4.7.2
<b>emaj_start_group</b>	group TEXT [ mark TEXT ] [ reset.log BOOLEAN ]	#.tables.and.seq INT		§ 4.2.4
<b>emaj_start_groups</b>	groups.array TEXT[] [ mark TEXT ] [ reset.log BOOLEAN ]	#.tables.and.seq INT		§ 4.2.4
<b>emaj_stop_group</b>	group TEXT [mark TEXT]	#.tables.and.seq INT		§ 4.2.8
<b>emaj_stop_groups</b>	groups.array TEXT[] [mark TEXT]	#.tables.and.seq INT		§ 4.2.8
<b>emaj_unprotect_group</b>	group TEXT	0/1 INT		§ 4.4.3
<b>emaj_unprotect_mark_group</b>	group TEXT mark TEXT	0/1 INT		§ 4.5.6
<b>emaj_verify_all</b>	-	Setof TEXT	Yes	§ 4.8.1

## 8.2 E-MAJ DISTRIBUTION CONTENT

Once installed (see §3), an E-Maj version contents the following files.

sql/emaj--3.1.0.sql	installation script of the extension
sql/emaj--3.0.1--3.1.0.sql	extension upgrade script from 3.0.0 to 3.1.0
sql/emaj--2.3.1--3.0.0.sql	extension upgrade script from 2.3.1 to 3.0.0
sql/emaj--2.3.0--2.3.1.sql	extension upgrade script from 2.3.0 to 2.3.1
sql/emaj--2.2.3--2.3.0.sql	extension upgrade script from 2.2.3 to 2.3.0
sql/emaj--2.2.2--2.2.3.sql	extension upgrade script from 2.2.2 to 2.2.3
sql/emaj--2.2.1--2.2.2.sql	extension upgrade script from 2.2.1 to 2.2.2
sql/emaj--2.2.0--2.2.1.sql	extension upgrade script from 2.2.0 to 2.2.1
sql/emaj--2.1.0--2.2.0.sql	extension upgrade script from 2.1.0 to 2.2.0
sql/emaj--2.0.1--2.1.0.sql	extension upgrade script from 2.0.1 to 2.1.0
sql/emaj--2.0.0--2.0.1.sql	extension upgrade script from 2.0.0 to 2.0.1
sql/emaj--1.3.1--2.0.0.sql	extension upgrade script from 1.3.1 to 2.0.0
sql/emaj--unpackaged--1.3.1.sql	script that transforms an existing 1.3.1 version into extension
sql/emaj-1.3.0-to-1.3.1.sql	psql script that upgrades a 1.3.0 version
sql/emaj-1.2.0-to-1.3.0.sql	psql script that upgrades a 1.2.0 version
sql/emaj-1.1.0-to-1.2.0.sql	psql script that upgrades a 1.1.0 version
sql/emaj-1.0.2-to-1.1.0.sql	psql script that upgrades a 1.0.2 version
sql/emaj-1.0.1-to-1.0.2.sql	psql script that upgrades a 1.0.1 version
sql/emaj-1.0.0-to-1.0.1.sql	psql script that upgrades a 1.0.0 version
sql/emaj-0.11.1-to-1.0.0.sql	psql script that upgrades a 0.11.1 version
sql/emaj-0.11.0-to-0.11.1.sql	psql script that upgrades a 0.11.0 version
sql/emaj-3.1.0.sql	psql script that creates the extension
sql/emaj_demo.sql	psql E-Maj demonstration script
sql/emaj_prepare_parallel_rollback_test.sql	psql test script for parallel rollbacks
sql/emaj_uninstall.sql	psql script to uninstall the E-Maj components
README.md	reduced extension's documentation
CHANGES.md	change log
LICENSE	information about E-Maj license
AUTHORS.md	who are the authors
META.json	technical data for PGXN
emaj.control	extension control file used by the integrated extensions management
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>_pres_en.odp	English version of the E-Maj presentation
doc/Emaj.<version>_pres_fr.odp	French version of the E-Maj presentation
doc/Emaj.<version>_pres_en.pdf	English version of the E-Maj presentation (pdf version)
doc/Emaj.<version>_pres_fr.pdf	French version of the E-Maj presentation (pdf version)
client/emajParallelRollback.php	php tool for parallel rollback
client/emajParallelRollback.pl	perl tool for parallel rollback
client/emajRollbackMonitor.php	php tool for rollbacks monitoring
client/emajRollbackMonitor.pl	perl tool for rollbacks monitoring

## 8.3 POSTGRESQL AND E-MAJ VERSIONS COMPATIBILITY MATRIX

---

PostgreSQL versions		E-Maj versions	
Min	Max	Min	Date
8.2	8.4	0.4.0	01/2010
8.2	9.0	0.8.0	10/2010
8.2	9.1	0.10.0	11/2011
8.2	9.2	0.11.1	07/2012
8.3	9.3	1.1.0	10/2013
8.3	9.5	1.2.0	01/2016
8.3	9.6	1.3.1	09/2016
9.1	9.6	2.0.0	11/2016
9.1	10	2.1.0	08/2017
9.2	10	2.3.0	07/2018
9.2	11	2.3.1	09/2018
9.5	11	3.0.0	03/2019
9.5	12	3.1.0	06/2019