

# Getting on a hook or PostgreSQL extensibility

Alexey Kondratov

Postgres Professional

PostgreSQL @ FOSDEM'21, February 6-7

# PostgreSQL extensibility

- Custom types, operators.
- Access methods.
- PL/pgSQL scripting language.
- Functions, triggers, extensions and so on.
- ...



A lot of info in the docs [1].

# PostgreSQL extensibility

- Custom types, operators.
- Access methods.
- PL/pgSQL scripting language.
- Functions, triggers, extensions and so on.
- ...
- **Hooks** (*and callbacks*).



A lot of info in the docs [1].

# What is a hook?

- Function or more precisely **a global pointer to a function.**
- Being defined it will be called by PostgreSQL at some specific moment.
- Scattered all over the PostgreSQL core.
- Extensions (shared libraries) can set these hooks to peek into the PostgreSQL internal state.

# Hooks: pointer

```
136  * We provide a function hook variable that lets loadable plugins
137  * get control when ExecutorStart is called.  Such a plugin would
138  * normally call standard_ExecutorStart().
139  *
140  * -----
141  */
142  void
143  ExecutorStart(QueryDesc *queryDesc, int eflags)
144  {
145      if (ExecutorStart_hook)
146          (*ExecutorStart_hook) (queryDesc, eflags);
147      else
148          standard_ExecutorStart(queryDesc, eflags);
149  }
```

execMain.c

Executed if defined

# Hooks: installation

```
473  /*
474  * Install hooks.
475  */
476  prev_shmem_startup_hook = shmem_startup_hook;
477  shmem_startup_hook = pgss_shmem_startup;
478  prev_post_parse_analyze_hook = post_parse_analyze_hook;
479  post_parse_analyze_hook = pgss_post_parse_analyze;
480  prev_planner_hook = planner_hook;
481  planner_hook = pgss_planner;
482  prev_ExecutorStart = ExecutorStart_hook;
483  ExecutorStart_hook = pgss_ExecutorStart;
```

pg\_stat\_statements.c: \_PG\_init()

- 1) Remember previously defined value
- 2) Register your own function

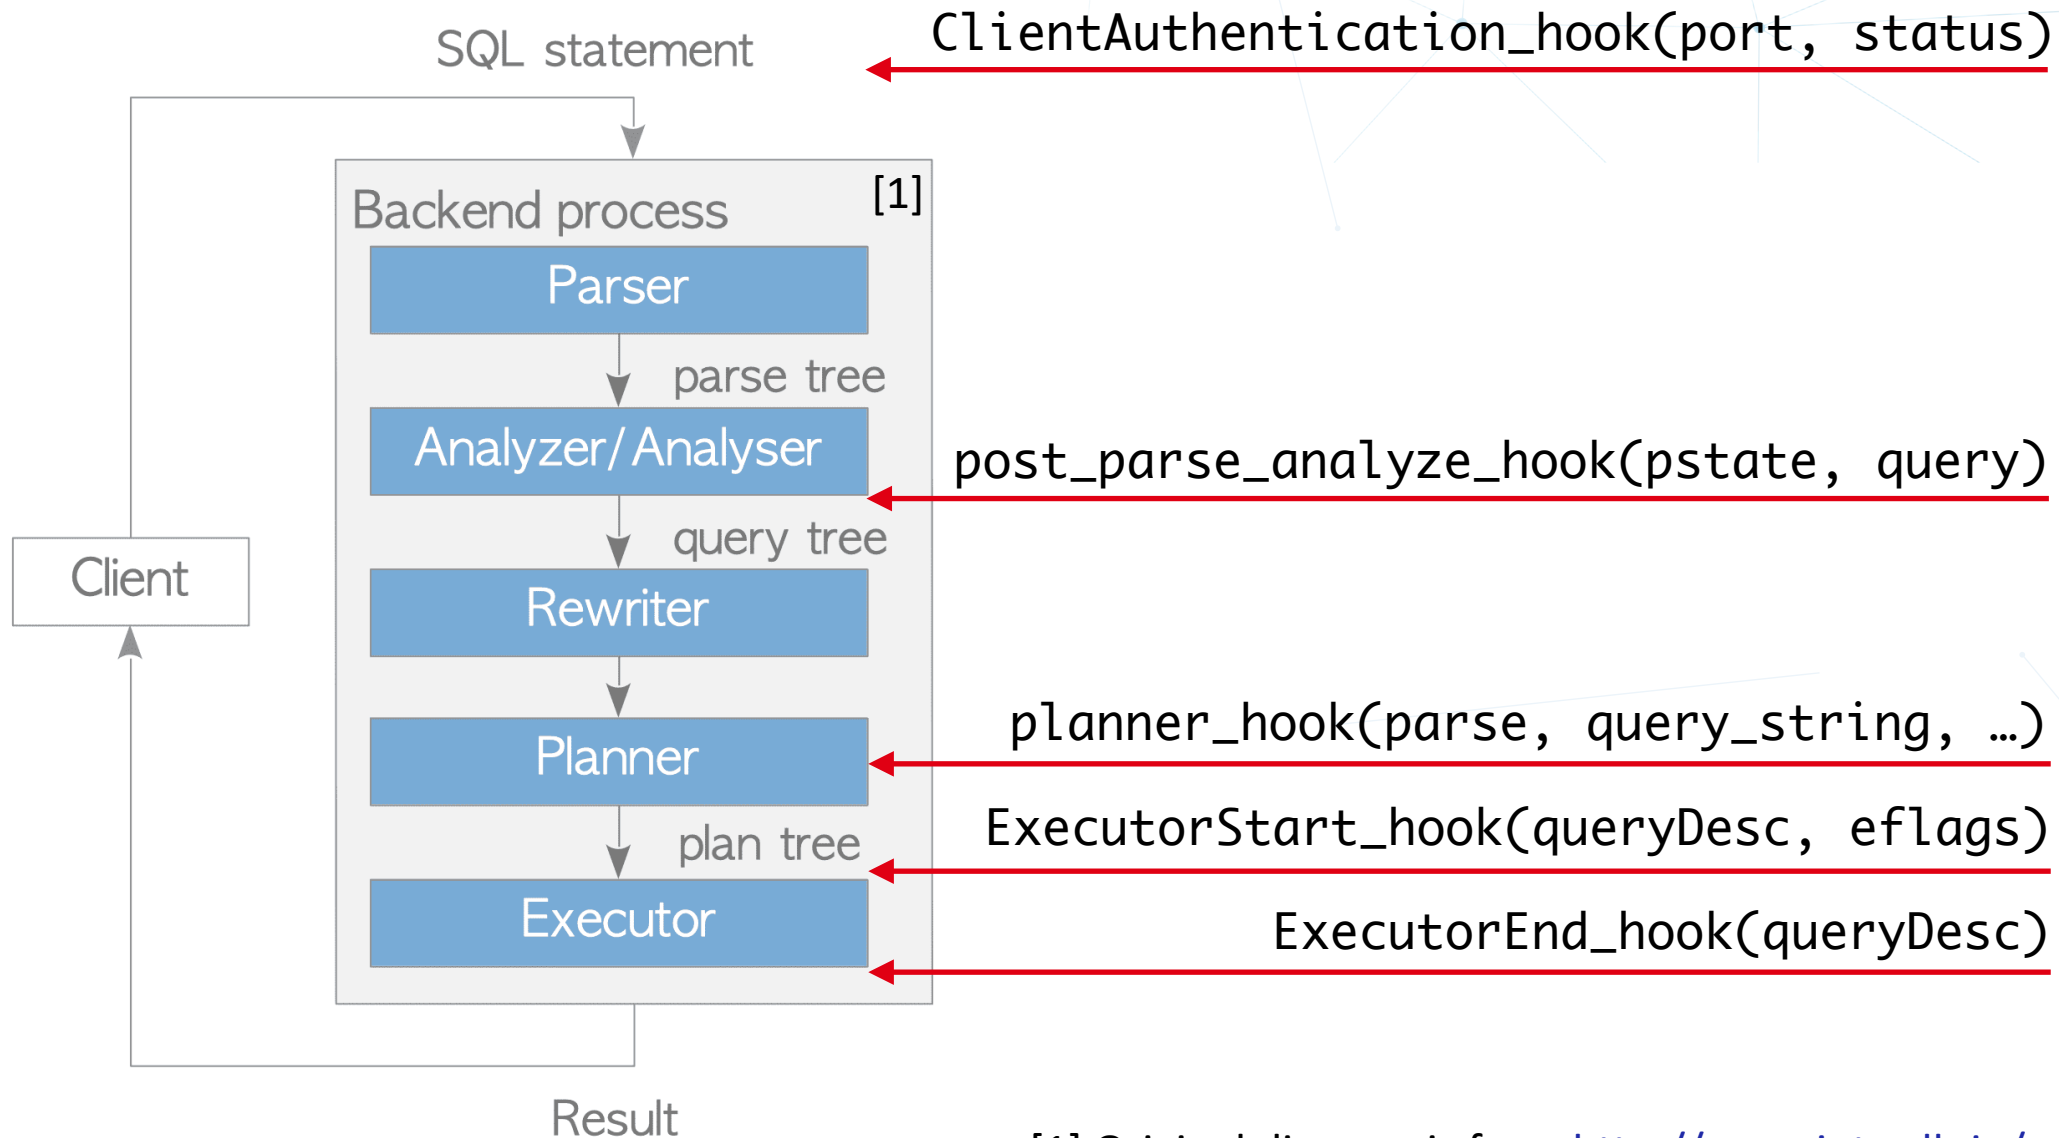
# Hooks: design

pg\_stat\_statements.c

```
998  /*
999  * ExecutorStart hook: start up tracking if needed
1000  */
1001  static void
1002  pgss_ExecutorStart(QueryDesc *queryDesc, int eflags)
1003  {
1004      if (prev_ExecutorStart)
1005          prev_ExecutorStart(queryDesc, eflags);
1006      else
1007          standard_ExecutorStart(queryDesc, eflags);
```

Do not forget to call your predecessor!

# Hooks: overview





# Hooks: unofficial documentation

- GitHub repo: <https://github.com/AmatanHead/psql-hooks>
  - Lists hook arguments.
  - Has text description.
- pgPedia: <https://pgpedia.info/h/hooks.html>
  - There is an interesting change history with commit reference per hook.
- A bit outdated [Guillaume Lelarge's slides](#) from [PGCon 2012](#).

# What is a callback?

- Very similar to the hooks.
- But initially designed to be set by multiple users.
- Usually installed by Register\*Callback() setter functions: RegisterXactCallback(), RegisterSubXactCallback(), RegisterExprContextCallback(), etc.
- Yet, there are others like: before\_shmem\_exit(), on\_shmem\_exit().
- Mostly for internal usage.

# Callbacks: registration

postgres\_fdw/connection.c

```
164
165      /*
166       * Register some callback functions that manage connection cleanup.
167       * This should be done just once in each backend.
168       */
169      RegisterXactCallback(pgfdw_xact_callback, NULL);
170      RegisterSubXactCallback(pgfdw_subxact_callback, NULL);
171      CacheRegisterSyscacheCallback(FOREIGNSERVEROID,
172                                   pgfdw_inval_callback, (Datum) 0);
173      CacheRegisterSyscacheCallback(USERMAPPINGOID,
174                                   pgfdw_inval_callback, (Datum) 0);
```



Run setter function to register your own callback

# Callbacks: setter function

```
3535 void
3536 RegisterXactCallback(XactCallback callback, void *arg)
3537 {
3538     XactCallbackItem *item;
3539
3540     item = (XactCallbackItem *)
3541         MemoryContextAlloc(TopMemoryContext, sizeof(XactCallbackItem));
3542     item->callback = callback;
3543     item->arg = arg;
3544     item->next = Xact_callbacks;
3545     Xact_callbacks = item;
3546 }
```

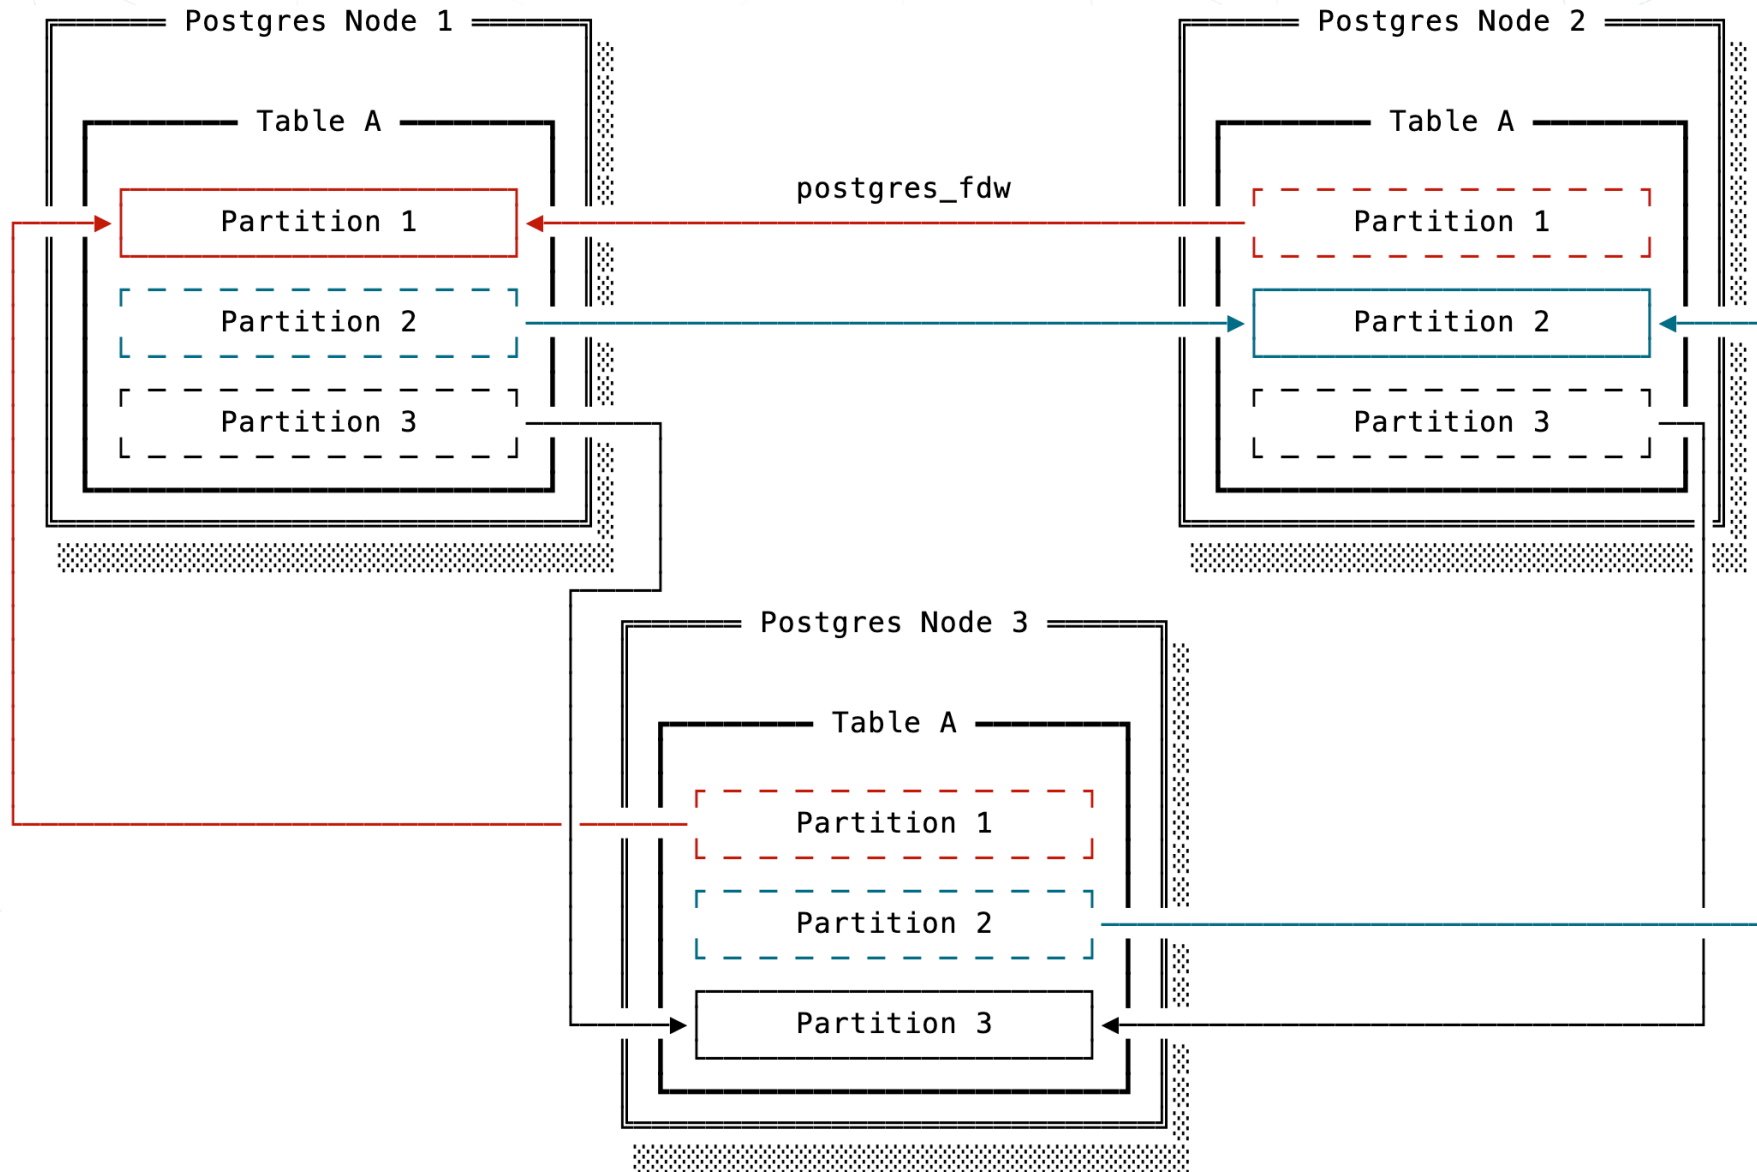
xact.c

Keeps a list of registered callbacks



Example time

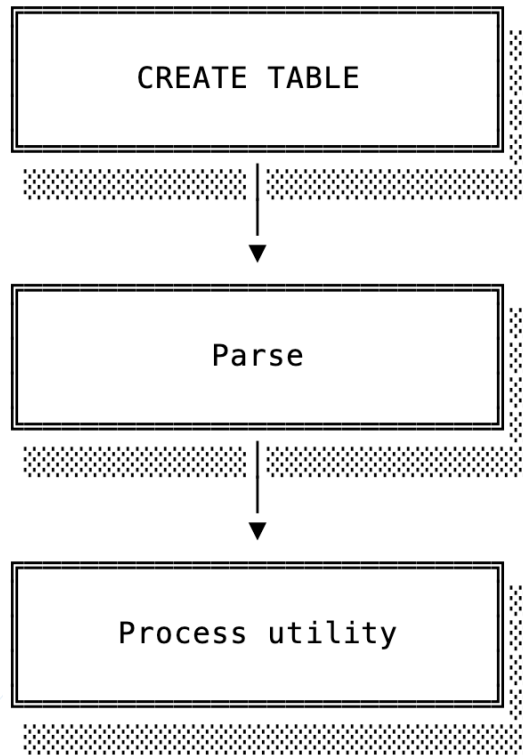
# Distributed PostgreSQL



# Distributed DDL

- **Broadcast** specific (or all) DDL across a number of PostgreSQL nodes.
- Create distributed (sharded / partitioned) tables with familiar interface → **extend CREATE TABLE statement syntax.**
- This operation should be **atomic**, i.e. either committed or aborted on all PostgreSQL instances → **use two-phase commit (2PC).**
- Do everything from the extension → **no core modifications!**

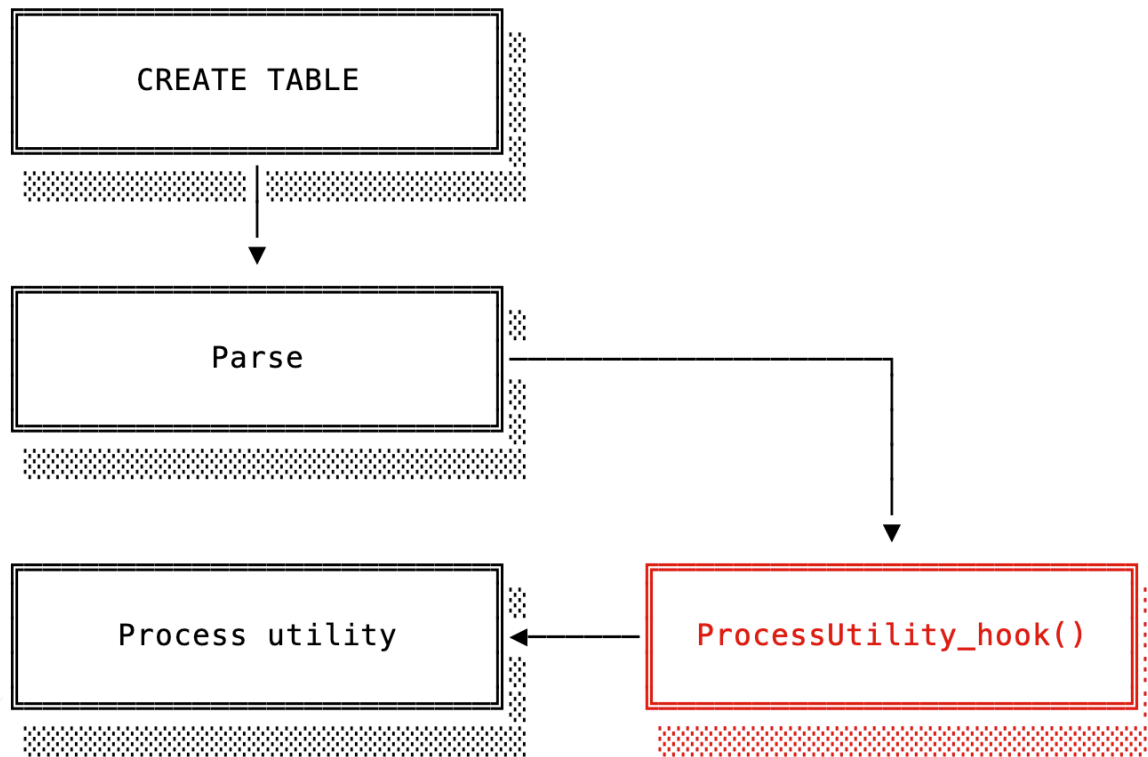
# Standard DDL processing



- Get query from the client.
- Parse and plan it.
- Pass it to the `standard_ProcessUtility()`.



# Distributed DDL: broadcast




Utility hook receives:

- Raw text of the statement.
- Planned statement.
- **So it can decide whether to send this DDL to other servers or not [1].**

# Distributed DDL: syntax extension


```
1
2 CREATE TABLE users (
3   id      int not null,
4   name    text
5 ) WITH (distributed_by = 'id',
6        num_parts = 12,
7        colocate_with = 'companies');
8
```



We would like to add some additional parameters to CREATE TABLE syntax (e.g. **number of partitions**, **partitioning column name**).

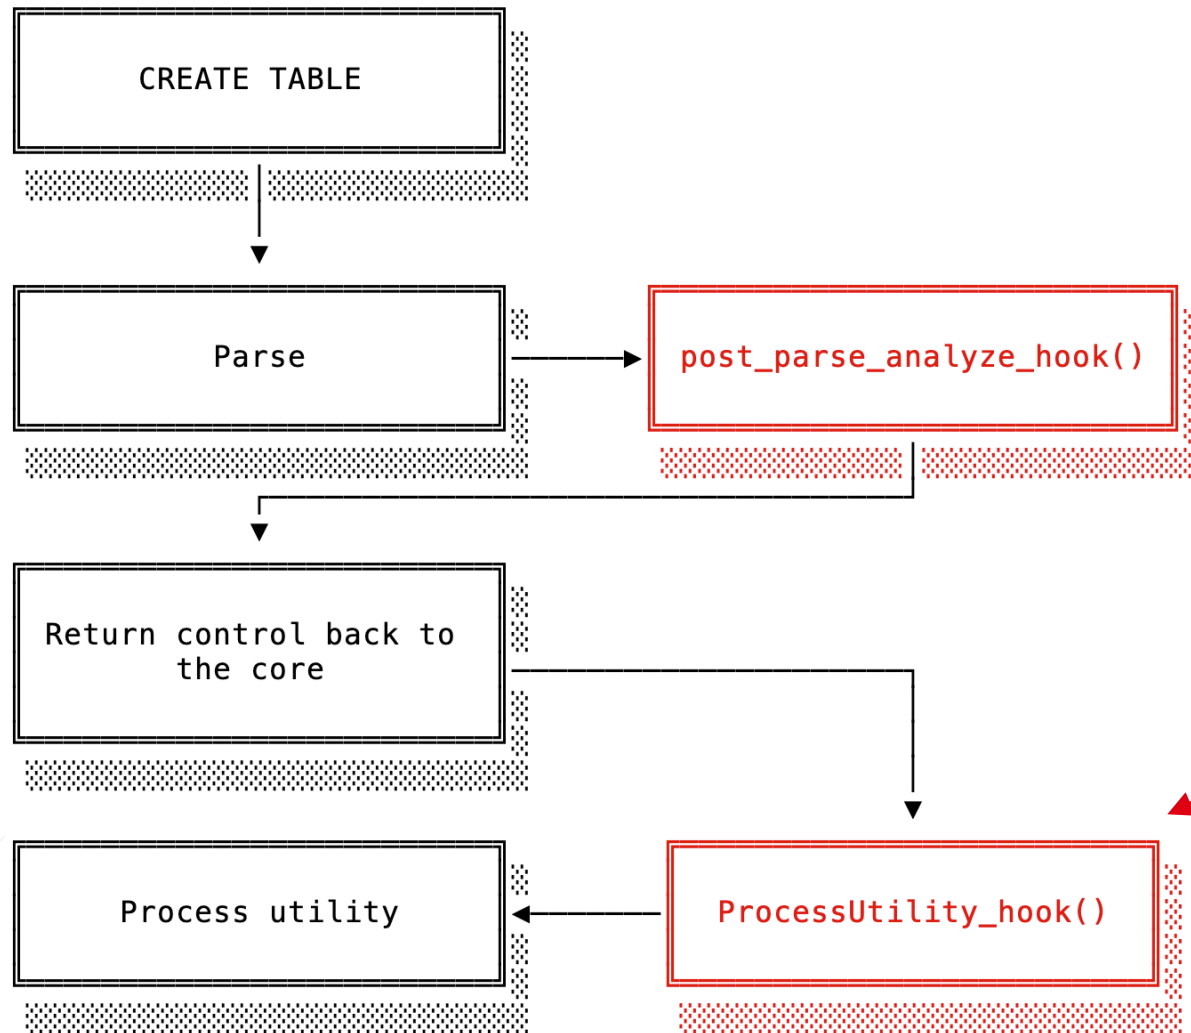
# Distributed DDL: syntax extension

```
pid:13864 [local]:5432 alexk@postgres=# CREATE TABLE users (  
  id      int not null,  
  name    text  
) WITH (distributed_by = 'id',  
        num_parts = 12,  
        colocate_with = 'companies');  
ERROR: unrecognized parameter "distributed_by"
```



Luckily, not a ‘syntax error’, so **parameters are not processed by the parser itself!**

# Distributed DDL: syntax extension



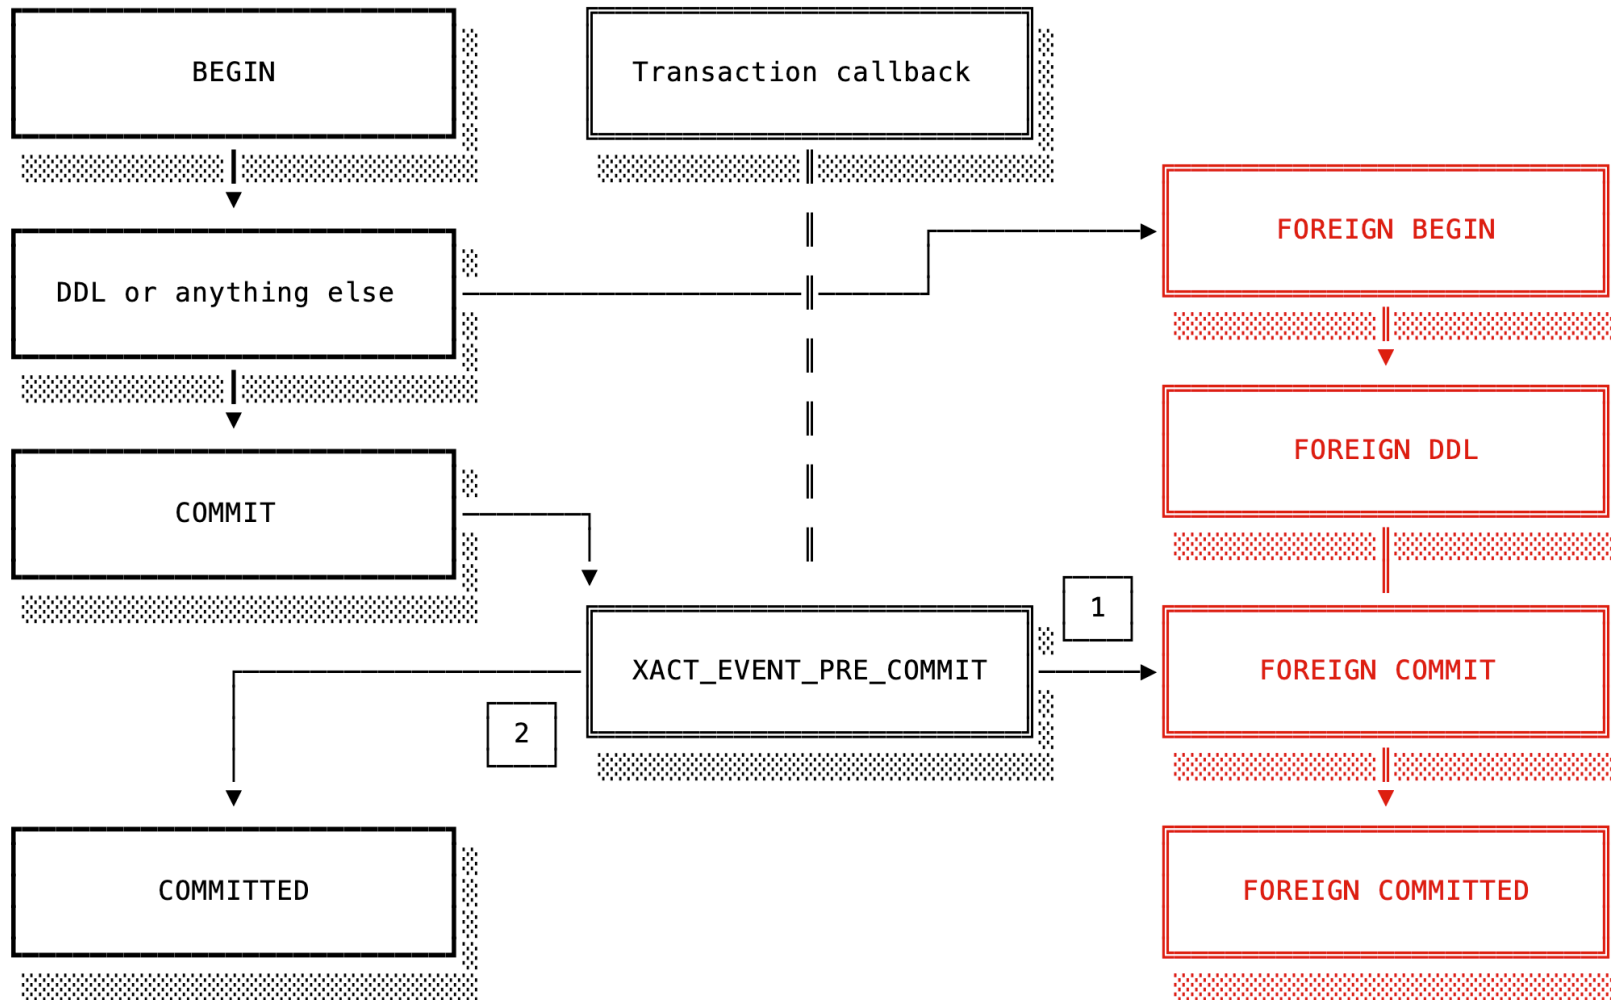
1. Notice, remember and **remove** known additional parameters.

2. Process statement taking into account the specified parameters (i.e. **add partitioning info, create partitions as well, do broadcast**).

# Distributed DDL: atomicity

- Without 2PC, transaction might end up **COMMITTED** on some nodes and **ABORTED** on others.
- 2PC introduces an intermediate state — **PREPARED**.
- PostgreSQL already has a 2PC infrastructure.

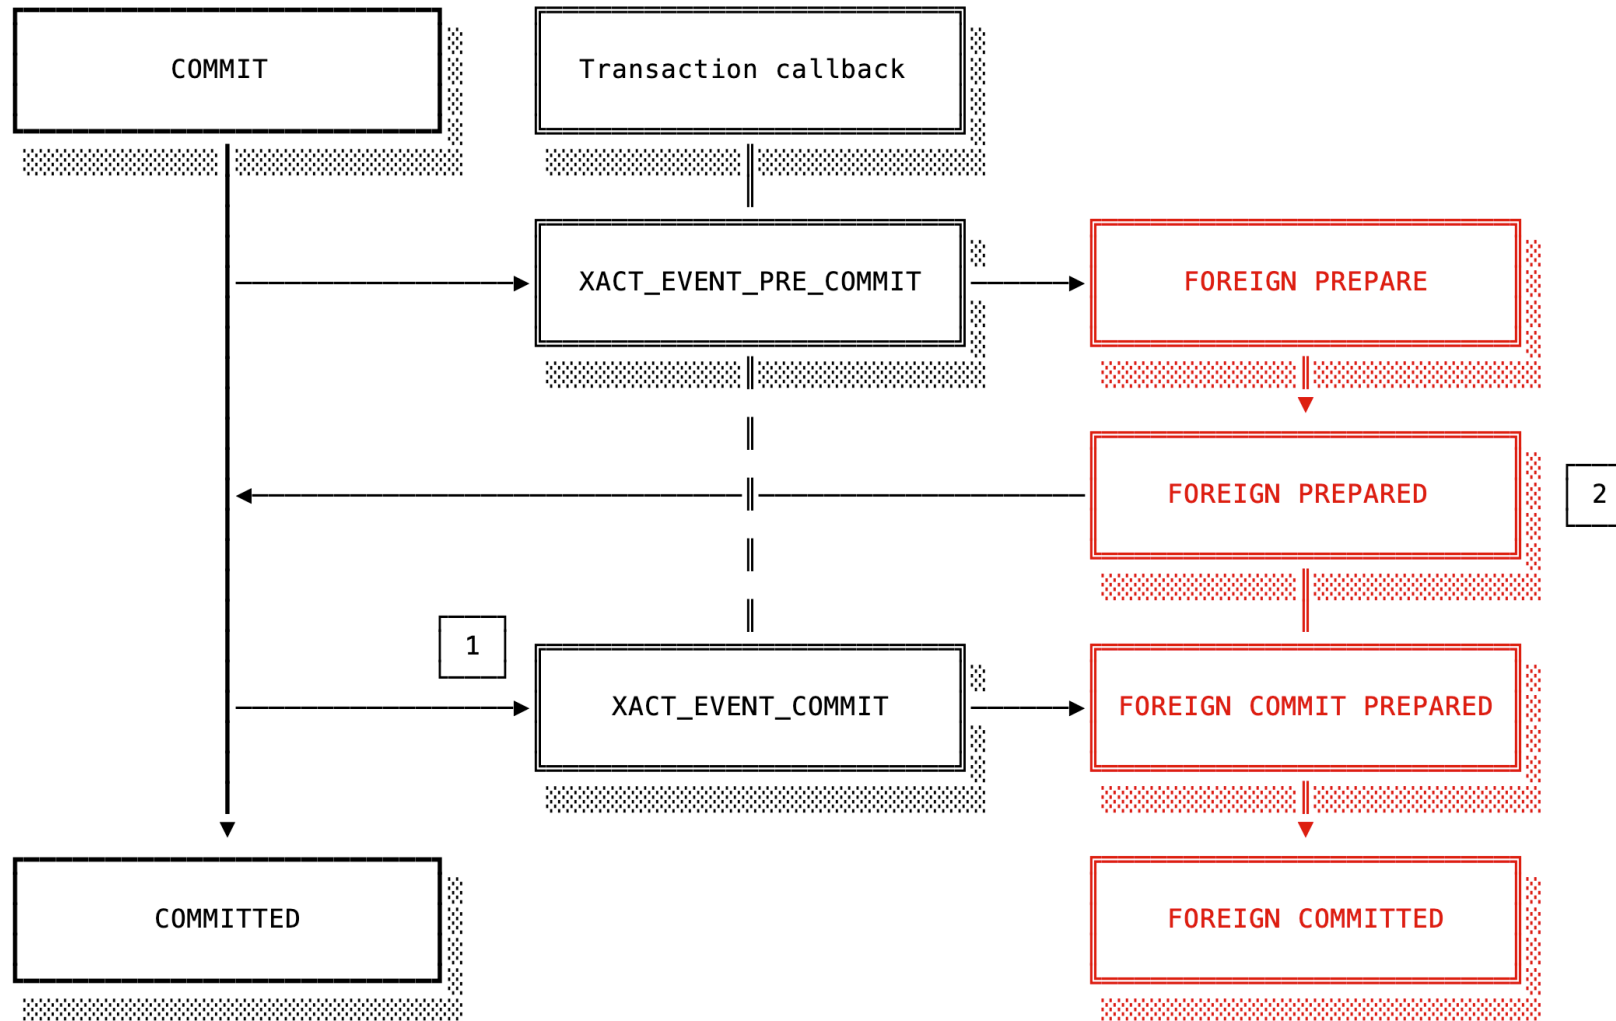
# Distributed DDL: transaction



Transaction (xact) callback  
is used by postgres\_fdw to:

1. Commit all foreign transactions first.
2. Proceed to local commit.

# Distributed DDL: 2PC



At stage (1) it is too late to abort local transaction and if we will fail to commit all remote xacts, then some of them may be left in the **PREPARED** state (2). In this case some additional process (resolver) have to either commit them or abort based on the coordinator state.

# Feedback

If you have any questions or comments:

- [kondratov.aleksey@gmail.com](mailto:kondratov.aleksey@gmail.com)
- [github.com/ololobus](https://github.com/ololobus)
- [twitter.com/ololobuss](https://twitter.com/ololobuss)

Thank you!