



Mnesia

Copyright © 1997-2020 Ericsson AB. All Rights Reserved.
Mnesia 4.17
June 21, 2020

Copyright © 1997-2020 Ericsson AB. All Rights Reserved.

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0> Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License. Ericsson AB. All Rights Reserved..

June 21, 2020

1.3 Getting Started

```
unix> erl -mnesia dir '"/tmp/funky"'
Erlang (BEAM) emulator version 4.9

Eshell V4.9 (abort with ^G)
1>
1> mnesia:create_schema([node()]).
ok
2> mnesia:start().
ok
3> mnesia:create_table(funky, []).
{atomic,ok}
4> mnesia:info().
---> Processes holding locks <---
---> Processes waiting for locks <---
---> Pending (remote) transactions <---
---> Active (local) transactions <---
---> Uncertain transactions <---
---> Active tables <---
funky          : with 0 records occupying 269 words of mem
schema         : with 2 records occupying 353 words of mem
===> System info in version "1.0", debug level = none <===
opt_disc. Directory "/tmp/funky" is used.
use fall-back at restart = false
running db nodes = [nonode@nohost]
stopped db nodes = []
remote          = []
ram_copies      = [funky]
disc_copies     = [schema]
disc_only_copies = []
[{nonode@nohost,disc_copies}] = [schema]
[{nonode@nohost,ram_copies}] = [funky]
1 transactions committed, 0 aborted, 0 restarted, 1 logged to disc
0 held locks, 0 in queue; 0 local transactions, 0 remote
0 transactions waits for other nodes: []
ok
```

In this example, the following actions are performed:

- **Step 1:** The Erlang system is started from the UNIX prompt with a flag `-mnesia dir '"/tmp/funky"'`, which indicates in which directory to store the data.
- **Step 2:** A new empty schema is initialized on the local node by evaluating `mnesia:create_schema([node()])`. The schema contains information about the database in general. This is explained in detail later.
- **Step 3:** The DBMS is started by evaluating `mnesia:start()`.
- **Step 4:** A first table is created, called `funky`, by evaluating the expression `mnesia:create_table(funky, [])`. The table is given default properties.
- **Step 5:** `mnesia:info()` is evaluated to display information on the terminal about the status of the database.

1.3.2 Example

A *Mnesia* database is organized as a set of tables. Each table is populated with instances (Erlang records). A table has also a number of properties, such as location and persistence.

Database

This example shows how to create a database called *Company* and the relationships shown in the following diagram:



Figure 3.1: Company Entity-Relation Diagram

The database model is as follows:

- There are three entities: department, employee, and project.
- There are three relationships between these entities:
 - A department is managed by an employee, hence the `manager` relationship.
 - An employee works at a department, hence the `at_dep` relationship.
 - Each employee works on a number of projects, hence the `in_proj` relationship.

Defining Structure and Content

First the record definitions are entered into a text file named `company.hrl`. This file defines the following structure for the example database:

```

-record(employee, {emp_no,
                  name,
                  salary,
                  sex,
                  phone,
                  room_no}).

-record(dept, {id,
              name}).

-record(project, {name,
                 number}).

-record(manager, {emp,
                 dept}).

-record(at_dep, {emp,
                dept_id}).

-record(in_proj, {emp,
                 proj_name}).
  
```

The structure defines six tables in the database. In `Mnesia`, the function `mnesia:create_table(Name, ArgList)` creates tables. `Name` is the table name.

Note:

The current version of Mnesia does not require that the name of the table is the same as the record name, see Record Names versus Table Names..

For example, the table for employees is created with the function `mnesia:create_table(employee, [{attributes, record_info(fields, employee)}])`. The table name `employee` matches the name for records specified in `ArgList`. The expression `record_info(fields, RecordName)` is processed by the Erlang preprocessor and evaluates to a list containing the names of the different fields for a record.

Program

The following shell interaction starts Mnesia and initializes the schema for the Company database:

```
% erl -mnesia dir '/ldisc/scratch/Mnesia.Company'
Erlang (BEAM) emulator version 4.9

Eshell V4.9 (abort with ^G)
1> mnesia:create_schema([node()]).
ok
2> mnesia:start().
ok
```

The following program module creates and populates previously defined tables:

```
-include_lib("stdlib/include/qlc.hrl").
-include("company.hrl").

init() ->
    mnesia:create_table(employee,
        [{attributes, record_info(fields, employee)}]),
    mnesia:create_table(dept,
        [{attributes, record_info(fields, dept)}]),
    mnesia:create_table(project,
        [{attributes, record_info(fields, project)}]),
    mnesia:create_table(manager, [{type, bag},
        {attributes, record_info(fields, manager)}]),
    mnesia:create_table(at_dep,
        [{attributes, record_info(fields, at_dep)}]),
    mnesia:create_table(in_proj, [{type, bag},
        {attributes, record_info(fields, in_proj)}]).
```

Program Explained

The following commands and functions are used to initiate the Company database:

- `% erl -mnesia dir '/ldisc/scratch/Mnesia.Company'`. This is a UNIX command-line entry that starts the Erlang system. The flag `-mnesia dir Dir` specifies the location of the database directory. The system responds and waits for further input with the prompt `1>`.
- `mnesia:create_schema([node()])`. This function has the format `mnesia:create_schema(DiscNodeList)` and initiates a new schema. In this example, a non-distributed system using only one node is created. Schemas are fully explained in Define a Schema.
- `mnesia:start()`. This function starts Mnesia and is fully explained in Start Mnesia.

Continuing the dialogue with the Erlang shell produces the following:

If the startup procedure fails, the function `mnesia:start()` returns the cryptic tuple `{error, {shutdown, {mnesia_sup, start_link, [normal, []]}}}`. To get more information about the start failure, use command-line arguments `-boot start_sasl` as argument to the `erl` script.

1.4.4 Create Tables

The function `mnesia:create_table(Name, ArgList)` creates tables. When executing this function, it returns one of the following responses:

- `{atomic, ok}` if the function executes successfully
- `{aborted, Reason}` if the function fails

The function arguments are as follows:

- `Name` is the name of the table. It is usually the same name as the name of the records that constitute the table. For details, see `record_name`.


```
mnesia:dirty_write(Record) ->
    Tab = element(1, Record),
    mnesia:dirty_write(Tab, Record).

mnesia:dirty_delete({Tab, Key}) ->
    mnesia:dirty_delete(Tab, Key).

mnesia:dirty_delete_object(Record) ->
    Tab = element(1, Record),
    mnesia:dirty_delete_object(Tab, Record)

mnesia:dirty_update_counter({Tab, Key}, Incr) ->
    mnesia:dirty_update_counter(Tab, Key, Incr).

mnesia:dirty_read({Tab, Key}) ->
    Tab = element(1, Record),
    mnesia:dirty_read(Tab, Key).

mnesia:dirty_match_object(Pattern) ->
    Tab = element(1, Pattern),
    mnesia:dirty_match_object(Tab, Pattern).

mnesia:dirty_index_match_object(Pattern, Attr)
    Tab = element(1, Pattern),
    mnesia:dirty_index_match_object(Tab, Pattern, Attr).

mnesia:write(Record) ->
    Tab = element(1, Record),
    mnesia:write(Tab, Record, write).

mnesia:s_write(Record) ->
    Tab = element(1, Record),
    mnesia:write(Tab, Record, sticky_write).

mnesia:delete({Tab, Key}) ->
    mnesia:delete(Tab, Key, write).

mnesia:s_delete({Tab, Key}) ->
    mnesia:delete(Tab, Key, sticky_write).

mnesia:delete_object(Record) ->
    Tab = element(1, Record),
    mnesia:delete_object(Tab, Record, write).

mnesia:s_delete_object(Record) ->
    Tab = element(1, Record),
    mnesia:delete_object(Tab, Record, sticky_write).

mnesia:read({Tab, Key}) ->
    mnesia:read(Tab, Key, read).

mnesia:wread({Tab, Key}) ->
    mnesia:read(Tab, Key, write).

mnesia:match_object(Pattern) ->
    Tab = element(1, Pattern),
    mnesia:match_object(Tab, Pattern, read).

mnesia:index_match_object(Pattern, Attr) ->
    Tab = element(1, Pattern),
    mnesia:index_match_object(Tab, Pattern, Attr, read).
```


1.6 Miscellaneous Mnesia Features

each fragment evenly over all the nodes in the node pool. Hopefully all nodes end up with the same number of replicas. `node_pool` defaults to the return value from the function `mnesia:system_info(db_nodes)`.

`{n_ram_copies, Int}`

Regulates how many `ram_copies` replicas that each fragment is to have. This property can explicitly be set at table creation. Defaults is 0, but if `n_disc_copies` and `n_disc_only_copies` also are 0, `n_ram_copies` defaults to 1.

`{n_disc_copies, Int}`

Regulates how many `disc_copies` replicas that each fragment is to have. This property can explicitly be set at table creation. Default is 0.

`{n_disc_only_copies, Int}`

Regulates how many `disc_only_copies` replicas that each fragment is to have. This property can explicitly be set at table creation. Defaults is 0.

`{foreign_key, ForeignKey}`

`ForeignKey` can either be the atom `undefined` or the tuple `{ForeignTab, Attr}`, where `Attr` denotes an attribute that is to be interpreted as a key in another fragmented table named `ForeignTab`. Mnesia ensures that the number of fragments in this table and in the foreign table are always the same.

When fragments are added or deleted, Mnesia automatically propagates the operation to all fragmented tables that have a foreign key referring to this table. Instead of using the record key to determine which fragment to access, the value of field `Attr` is used. This feature makes it possible to colocate records automatically in different tables to the same node. `foreign_key` defaults to `undefined`. However, if the foreign key is set to something else, it causes the default values of the other fragmentation properties to be the same values as the actual fragmentation properties of the foreign table.

`{hash_module, Atom}`

Enables definition of an alternative hashing scheme. The module must implement the `mnesia_frag_hash` callback behavior. This property can explicitly be set at table creation. Default is `mnesia_frag_hash`.

`{hash_state, Term}`

Enables a table-specific parameterization of a generic hash module. This property can explicitly be set at table creation. Default is `undefined`.

```

Eshell V4.7.3.3 (abort with ^G)
(a@sam)1> mnesia:start().
ok
(a@sam)2> PrimProps = [{n_fragments, 7}, {node_pool, [node()]}].
[{n_fragments,7},{node_pool,[a@sam]]}
(a@sam)3> mnesia:create_table(prim_dict,
                             [{frag_properties, PrimProps},
                              {attributes,[prim_key,prim_val]})}.
{atomic,ok}
(a@sam)4> SecProps = [{foreign_key, {prim_dict, sec_val}}].
[{foreign_key,{prim_dict,sec_val}}]
(a@sam)5> mnesia:create_table(sec_dict,
                             [{frag_properties, SecProps},
                              {attributes, [sec_key, sec_val]})}.
(a@sam)5>
{atomic,ok}
(a@sam)6> Write = fun(Rec) -> mnesia:write(Rec) end.
#Fun<erl_eval>
(a@sam)7> PrimKey = 11.
11
(a@sam)8> SecKey = 42.
42
(a@sam)9> mnesia:activity(sync_dirty, Write,
                          [{prim_dict, PrimKey, -11}], mnesia_frag).
ok
(a@sam)10> mnesia:activity(sync_dirty, Write,
                           [{sec_dict, SecKey, PrimKey}], mnesia_frag).
ok
(a@sam)11> mnesia:change_table_frag(prim_dict, {add_frag, [node()]}).
{atomic,ok}
(a@sam)12> SecRead = fun(PrimKey, SecKey) ->
               mnesia:read({sec_dict, PrimKey}, SecKey, read) end.
#Fun<erl_eval>
(a@sam)13> mnesia:activity(transaction, SecRead,
                           [PrimKey, SecKey], mnesia_frag).
[{sec_dict,42,11}]
(a@sam)14> Info = fun(Tab, Item) -> mnesia:table_info(Tab, Item) end.
#Fun<erl_eval>
(a@sam)15> mnesia:activity(sync_dirty, Info,
                           [prim_dict, frag_size], mnesia_frag).
[{prim_dict,0},
 {prim_dict_frag2,0},
 {prim_dict_frag3,0},
 {prim_dict_frag4,1},
 {prim_dict_frag5,0},
 {prim_dict_frag6,0},
 {prim_dict_frag7,0},
 {prim_dict_frag8,0}]
(a@sam)16> mnesia:activity(sync_dirty, Info,
                           [sec_dict, frag_size], mnesia_frag).
[{sec_dict,0},
 {sec_dict_frag2,0},
 {sec_dict_frag3,0},
 {sec_dict_frag4,1},
 {sec_dict_frag5,0},
 {sec_dict_frag6,0},
 {sec_dict_frag7,0},
 {sec_dict_frag8,0}]
(a@sam)17>

```

Management of Fragmented Tables

The function `mnesia:change_table_frag(Tab, Change)` is intended to be used for reconfiguration of fragmented tables. Argument `Change` is to have one of the following values:

- `checkpoints`. Returns the names of the currently active checkpoints, which involve this table on this node.
- `cookie`. Returns a table cookie, which is a unique system-generated identifier for the table. The cookie is used internally to ensure that two different table definitions using the same table name cannot accidentally be intermixed. The cookie is generated when the table is created initially.
- `disc_copies`. Returns the nodes where a `disc_copy` of the table resides according to the schema.
- `disc_only_copies`. Returns the nodes where a `disc_only_copy` of the table resides according to the schema.
- `index`. Returns the list of index position integers for the table.
- `load_node`. Returns the name of the node that Mnesia loaded the table from. The structure of the returned value is unspecified, but can be useful for debugging purposes.
- `load_order`. Returns the load order priority of the table. It is an integer and defaults to 0 (zero).
- `load_reason`. Returns the reason of why Mnesia decided to load the table. The structure of the returned value is unspecified, but can be useful for debugging purposes.
- `local_content`. Returns `true` or `false` to indicate if the table is configured to have locally unique content on each node.
- `master_nodes`. Returns the master nodes of a table.
- `memory`. Returns for `ram_copies` and `disc_copies` tables the number of words allocated in memory to the table on this node. For `disc_only_copies` tables the number of bytes stored on disc is returned.
- `ram_copies`. Returns the nodes where a `ram_copy` of the table resides according to the schema.
- `record_name`. Returns the record name, common for all records in the table.
- `size`. Returns the number of records inserted in the table.
- `snmp`. Returns the SNMP struct. `[]` means that the table currently has no SNMP properties.
- `storage_type`. Returns the local storage type of the table. It can be `disc_copies`, `ram_copies`, `disc_only_copies`, or the atom `unknown`. `unknown` is returned for all tables that only reside remotely.
- `subscribers`. Returns a list of local processes currently subscribing to local table events that involve this table on this node.
- `type`. Returns the table type, which is `bag`, `set`, or `ordered_set`.
- `user_properties`. Returns the user-associated table properties of the table. It is a list of the stored property records.
- `version`. Returns the current version of the table definition. The table version is incremented when the table definition is changed. The table definition can be incremented directly when it has been changed in a schema transaction, or when a committed table definition is merged with table definitions from other nodes during startup.
- `where_to_read`. Returns the node where the table can be read. If value `nowhere` is returned, either the table is not loaded or it resides at a remote node that is not running.
- `where_to_write`. Returns a list of the nodes that currently hold an active replica of the table.
- `wild_pattern`. Returns a structure that can be given to the various match functions for a certain table. A record tuple is where all record fields have value `'_'`.

```
transaction(Fun) -> t_result(Res)
transaction(Fun, Retries) -> t_result(Res)
transaction(Fun, Args :: [Arg :: term()]) -> t_result(Res)
transaction(Fun, Args :: [Arg :: term()], Retries) ->
    t_result(Res)
```

Types:

```
Fun = fun(...) -> Res)
Retries = integer() >= 0 | infinity
```

Executes the functional object `Fun` with arguments `Args` as a transaction.

The code that executes inside the transaction can consist of a series of table manipulation functions. If something goes wrong inside the transaction as a result of a user error or a certain table not being available, the entire transaction is terminated and the function `transaction/1` returns the tuple `{aborted, Reason}`.

If all is going well, `{atomic, ResultOfFun}` is returned, where `ResultOfFun` is the value of the last expression in `Fun`.

A function that adds a family to the database can be written as follows if there is a structure `{family, Father, Mother, ChildrenList}`:

```
add_family({family, F, M, Children}) ->
  ChildOids = lists:map(fun oid/1, Children),
  Trans = fun() ->
    mnesia:write(F#person{children = ChildOids},
    mnesia:write(M#person{children = ChildOids},
    Write = fun(Child) -> mnesia:write(Child) end,
    lists:foreach(Write, Children)
  end,
  mnesia:transaction(Trans).

oid(Rec) -> {element(1, Rec), element(2, Rec)}.
```

This code adds a set of people to the database. Running this code within one transaction ensures that either the whole family is added to the database, or the whole transaction terminates. For example, if the last child is badly formatted, or the executing process terminates because of an 'EXIT' signal while executing the family code, the transaction terminates. Thus, the situation where half a family is added can never occur.

It is also useful to update the database within a transaction if several processes concurrently update the same records. For example, the function `raise(Name, Amount)`, which adds `Amount` to the salary field of a person, is to be implemented as follows:

```
raise(Name, Amount) ->
  mnesia:transaction(fun() ->
    case mnesia:wread({person, Name}) of
      [P] ->
        Salary = Amount + P#person.salary,
        P2 = P#person{salary = Salary},
        mnesia:write(P2);
      _ ->
        mnesia:abort("No such person")
    end
  end).
```

When this function executes within a transaction, several processes running on different nodes can concurrently execute the function `raise/2` without interfering with each other.

Since Mnesia detects deadlocks, a transaction can be restarted any number of times. This function attempts a restart as specified in `Retries`. `Retries` must be an integer greater than 0 or the atom `infinity`. Default is `infinity`.

```
transform_table(Tab :: table(), Fun, NewA :: [Attr], RecName) ->
  t_result(ok)
```

Types:

```
RecName = Attr = atom()
Fun =
    fun((Record :: tuple()) -> Transformed :: tuple()) | ignore
```

Applies argument `Fun` to all records in the table. `Fun` is a function that takes a record of the old type and returns a transformed record of the new type. Argument `Fun` can also be the atom `ignore`, which indicates that only the metadata about the table is updated. Use of `ignore` is not recommended, but included as a possibility for the user to do an own transformation.

`NewAttributeList` and `NewRecordName` specify the attributes and the new record type of the converted table. Table name always remains unchanged. If `record_name` is changed, only the Mnesia functions that use table identifiers work, for example, `mnesia:write/3` works, but not `mnesia:write/1`.

```
transform_table(Tab :: table(), Fun, NewA :: [Attr]) ->
    t_result(ok)
```

Types:

```
Attr = atom()
Fun =
    fun((Record :: tuple()) -> Transformed :: tuple()) | ignore
```

Calls `mnesia:transform_table(Tab, Fun, NewAttributeList, RecName)`, where `RecName` is `mnesia:table_info(Tab, record_name)`.

```
traverse_backup(Src :: term(), Dest :: term(), Fun, Acc) ->
    {ok, Acc} | {error, Reason :: term()}
traverse_backup(Src :: term(),
    SrcMod :: module(),
    Dest :: term(),
    DestMod :: module(),
    Fun, Acc) ->
    {ok, Acc} | {error, Reason :: term()}
```

Types:

```
Fun = fun((Items, Acc) -> {Items, Acc})
```

Iterates over a backup, either to transform it into a new backup, or read it. The arguments are explained briefly here. For details, see the User's Guide.

- `SourceMod` and `TargetMod` are the names of the modules that actually access the backup media.
- `Source` and `Target` are opaque data used exclusively by modules `SourceMod` and `TargetMod` to initialize the backup media.
- `Acc` is an initial accumulator value.
- `Fun(BackupItems, Acc)` is applied to each item in the backup. The `Fun` must return a tuple `{BackupItems, NewAcc}`, where `BackupItems` is a list of valid backup items, and `NewAcc` is a new accumulator value. The returned backup items are written in the target backup.
- `LastAcc` is the last accumulator value. This is the last `NewAcc` value that was returned by `Fun`.

```
uninstall_fallback() -> result()
```

Calls the function `mnesia:uninstall_fallback([scope, global])`.

```
uninstall_fallback(Args) -> result()
```

Types:

```
Args = [{mnesia_dir, Dir :: string()}]
```

Deinstalls a fallback before it has been used to restore the database. This is normally a distributed operation that is either performed on all nodes with disc resident schema, or none. Uninstallation of fallbacks requires Erlang to be operational on all involved nodes, but it does not matter if Mnesia is running or not. Which nodes that are considered as disc-resident nodes is determined from the schema information in the local fallback.

Args is a list of the following tuples:

- {module, BackupMod}. For semantics, see `mnesia:install_fallback/2`.
- {scope, Scope}. For semantics, see `mnesia:install_fallback/2`.
- {mnesia_dir, AlternateDir}. For semantics, see `mnesia:install_fallback/2`.

```
unsubscribe(What) -> {ok, node()} | {error, Reason :: term()}
```

Types:

```
What = system | activity | {table, table()}, simple | detailed}
```

Stops sending events of type `EventCategory` to the caller.

Node is the local node.

```
wait_for_tables(Tabs :: [Tab :: table()], TMO :: timeout()) ->  
    result() | {timeout, [table()]}
```

Some applications need to wait for certain tables to be accessible to do useful work. `mnesia:wait_for_tables/2` either hangs until all tables in `TabList` are accessible, or until timeout is reached.

```
wread(Objid :: {Tab :: table(), Key :: term()}) -> [tuple()]
```

Calls the function `mnesia:read(Tab, Key, write)`.

```
write(Record :: tuple()) -> ok
```

Calls the function `mnesia:write(Tab, Record, write)`, where `Tab` is `element(1, Record)`.

```
write(Tab :: table(),  
      Record :: tuple(),  
      LockKind :: write_locks()) ->  
    ok
```

Writes record `Record` to table `Tab`.

The function returns `ok`, or terminates if an error occurs. For example, the transaction terminates if no `person` table exists.

The semantics of this function is context-sensitive. For details, see `mnesia:activity/4`. In transaction-context, it acquires a lock of type `LockKind`. The lock types `write` and `sticky_write` are supported.

```
write_lock_table(Tab :: table()) -> ok
```

Calls the function `mnesia:lock({table, Tab}, write)`.

Configuration Parameters

Mnesia reads the following application configuration parameters:

- `-mnesia access_module` Module. The name of the Mnesia activity access callback module. Default is `mnesia`.
- `-mnesia auto_repair` `true` | `false`. This flag controls if Mnesia automatically tries to repair files that have not been properly closed. Default is `true`.
- `-mnesia backup_module` Module. The name of the Mnesia backup callback module. Default is `mnesia_backup`.
- `-mnesia debug` Level. Controls the debug level of Mnesia. The possible values are as follows:

`none`

No trace outputs. This is the default.

`verbose`

Activates tracing of important debug events. These events generate `{mnesia_info, Format, Args}` system events. Processes can subscribe to these events with `mnesia:subscribe/1`. The events are always sent to the Mnesia event handler.

`debug`

Activates all events at the verbose level plus full trace of all debug events. These debug events generate `{mnesia_info, Format, Args}` system events. Processes can subscribe to these events with `mnesia:subscribe/1`. The events are always sent to the Mnesia event handler. On this debug level, the Mnesia event handler starts subscribing to updates in the schema table.

`trace`

Activates all events at the debug level. On this level, the Mnesia event handler starts subscribing to updates on all Mnesia tables. This level is intended only for debugging small toy systems, as many large events can be generated.

`false`

An alias for `none`.

`true`

An alias for `debug`.

- `-mnesia core_dir` Directory. The name of the directory where Mnesia core files is stored, or `false`. Setting it implies that also RAM-only nodes generate a core file if a crash occurs.
- `-mnesia dc_dump_limit` Number. Controls how often `disc_copies` tables are dumped from memory. Tables are dumped when `filesize(Log) > (filesize(Tab)/Dc_dump_limit)`. Lower values reduce CPU overhead but increase disk space and startup times. Default is 4.
- `-mnesia dir` Directory. The name of the directory where all Mnesia data is stored. The directory name must be unique for the current node. Two nodes must never share the the same Mnesia directory. The results are unpredictable.
- `-mnesia dump_disc_copies_at_startup` `true` | `false`. If set to `false`, this disables the dumping of `disc_copies` tables during startup while tables are being loaded. The default is `true`.
- `-mnesia dump_log_load_regulation` `true` | `false`. Controls if log dumps are to be performed as fast as possible, or if the dumper is to do its own load regulation. Default is `false`.

This feature is temporary and will be removed in a future release

- `-mnesia dump_log_update_in_place` `true` | `false`. Controls if log dumps are performed on a copy of the original data file, or if the log dump is performed on the original data file. Default is `true`
- `-mnesia dump_log_write_threshold` Max. Max is an integer that specifies the maximum number of writes allowed to the transaction log before a new dump of the log is performed. Default is 100 log writes.

- `-mnesia dump_log_time_threshold` `Max`. `Max` is an integer that specifies the dump log interval in milliseconds. Default is 3 minutes. If a dump has not been performed within `dump_log_time_threshold` milliseconds, a new dump is performed regardless of the number of writes performed.
- `-mnesia event_module` `Module`. The name of the Mnesia event handler callback module. Default is `mnesia_event`.
- `-mnesia extra_db_nodes` `Nodes` specifies a list of nodes, in addition to the ones found in the schema, with which Mnesia is also to establish contact. Default is `[]` (empty list).
- `-mnesia fallback_error_function` `{UserModule, UserFunc}`. Specifies a user-supplied callback function, which is called if a fallback is installed and Mnesia goes down on another node. Mnesia calls the function with one argument, the name of the dying node, for example, `UserModule:UserFunc(DyingNode)`. Mnesia must be restarted, otherwise the database can be inconsistent. The default behavior is to terminate Mnesia.
- `-mnesia max_wait_for_decision` `Timeout`. Specifies how long Mnesia waits for other nodes to share their knowledge about the outcome of an unclear transaction. By default, `Timeout` is set to the atom `infinity`. This implies that if Mnesia upon startup detects a "heavyweight transaction" whose outcome is unclear, the local Mnesia waits until Mnesia is started on some (in the worst case all) of the other nodes that were involved in the interrupted transaction. This is a rare situation, but if it occurs, Mnesia does not guess if the transaction on the other nodes was committed or terminated. Mnesia waits until it knows the outcome and then acts accordingly.

If `Timeout` is set to an integer value in milliseconds, Mnesia forces "heavyweight transactions" to be finished, even if the outcome of the transaction for the moment is unclear. After `Timeout` milliseconds, Mnesia commits or terminates the transaction and continues with the startup. This can lead to a situation where the transaction is committed on some nodes and terminated on other nodes. If the transaction is a schema transaction, the inconsistency can be fatal.

- `-mnesia no_table_loaders` `NUMBER`. Specifies the number of parallel table loaders during start. More loaders can be good if the network latency is high or if many tables contain few records. Default is 2.
- `-mnesia send_compressed` `Level`. Specifies the level of compression to be used when copying a table from the local node to another one. Default is 0.

`Level` must be an integer in the interval `[0, 9]`, where 0 means no compression and 9 means maximum compression. Before setting it to a non-zero value, ensure that the remote nodes understand this configuration.

- `-mnesia schema_location` `Loc`. Controls where Mnesia looks for its schema. Parameter `Loc` can be one of the following atoms:

`disc`

Mandatory disc. The schema is assumed to be located in the Mnesia directory. If the schema cannot be found, Mnesia refuses to start. This is the old behavior.

`ram`

Mandatory RAM. The schema resides in RAM only. At startup, a tiny new schema is generated. This default schema only contains the definition of the schema table and only resides on the local node. Since no other nodes are found in the default schema, configuration parameter `extra_db_nodes` must be used to let the node share its table definitions with other nodes.

Parameter `extra_db_nodes` can also be used on disc based nodes.

`opt_disc`

Optional disc. The schema can reside on disc or in RAM. If the schema is found on disc, Mnesia starts as a disc-based node and the storage type of the schema table is `disc_copies`. If no schema is found on disc, Mnesia starts as a disc-less node and the storage type of the schema table is `ram_copies`. Default value for the application parameter is `opt_disc`.

First, the SASL application parameters are checked, then the command-line flags are checked, and finally, the default value is chosen.

See Also

`application(3)`, `dets(3)`, `disk_log(3)`, `ets(3)`, `mnesia_registry(3)`, `qlc(3)`

mnesia_frag_hash

Erlang module

This module defines a callback behavior for user-defined hash functions of fragmented tables.

Which module that is selected to implement the `mnesia_frag_hash` behavior for a particular fragmented table is specified together with the other `frag_properties`. The `hash_module` defines the module name. The `hash_state` defines the initial hash state.

This module implements dynamic hashing, which is a kind of hashing that grows nicely when new fragments are added. It is well suited for scalable hash tables.

Exports

`init_state(Tab, State) -> NewState | abort(Reason)`

Types:

```
Tab = atom()
State = term()
NewState = term()
Reason = term()
```

Starts when a fragmented table is created with the function `mnesia:create_table/2` or when a normal (unfragmented) table is converted to be a fragmented table with `mnesia:change_table_frag/2`.

Notice that the function `add_frag/2` is started one time for each of the other fragments (except number 1) as a part of the table creation procedure.

State is the initial value of the `hash_state` `frag_property`. NewState is stored as `hash_state` among the other `frag_properties`.

`add_frag(State) -> {NewState, IterFrag, AdditionalLockFrag} | abort(Reason)`

Types:

```
State = term()
NewState = term()
IterFrag = [integer()]
AdditionalLockFrag = [integer()]
Reason = term()
```

To scale well, it is a good idea to ensure that the records are evenly distributed over all fragments, including the new one.

NewState is stored as `hash_state` among the other `frag_properties`.

As a part of the `add_frag` procedure, Mnesia iterates over all fragments corresponding to the `IterFrag` numbers and starts `key_to_frag_number(NewState, RecordKey)` for each record. If the new fragment differs from the old fragment, the record is moved to the new fragment.

As the `add_frag` procedure is a part of a schema transaction, Mnesia acquires write locks on the affected tables. That is, both the fragments corresponding to `IterFrag` and those corresponding to `AdditionalLockFrag`.

`del_frag(State) -> {NewState, IterFrag, AdditionalLockFrag} | abort(Reason)`

Types:

```
State = term()
NewState = term()
IterFrag = [integer()]
AdditionalLockFrag = [integer()]
Reason = term()
```

NewState is stored as hash_state among the other frag_properties.

As a part of the del_frag procedure, Mnesia iterates over all fragments corresponding to the IterFrag numbers and starts key_to_frag_number(NewState, RecordKey) for each record. If the new fragment differs from the old fragment, the record is moved to the new fragment.

Notice that all records in the last fragment must be moved to another fragment, as the entire fragment is deleted.

As the del_frag procedure is a part of a schema transaction, Mnesia acquires write locks on the affected tables. That is, both the fragments corresponding to IterFrag and those corresponding to AdditionalLockFrag.

key_to_frag_number(State, Key) -> FragNum | abort(Reason)

Types:

```
FragNum = integer()
Reason = term()
```

Starts whenever Mnesia needs to determine which fragment a certain record belongs to. It is typically started at read, write, and delete.

match_spec_to_frag_numbers(State, MatchSpec) -> FragNums | abort(Reason)

Types:

```
MatchSpec = ets_select_match_spec()
FragNums = [FragNum]
FragNum = integer()
Reason = term()
```

This function is called whenever Mnesia needs to determine which fragments that need to be searched for a MatchSpec. It is typically called by select and match_object.

See Also

mnesia(3)

mnesia_registry

Erlang module

This module is usually part of the `erl_interface` application, but is currently part of the Mnesia application.

This module is mainly intended for internal use within OTP, but it has two functions that are exported for public use.

On C-nodes, `erl_interface` has support for registry tables. These tables reside in RAM on the C-node, but can also be dumped into Mnesia tables. By default, the dumping of registry tables through `erl_interface` causes a corresponding Mnesia table to be created with `mnesia_registry:create_table/1`, if necessary.

Tables that are created with these functions can be administered as all other Mnesia tables. They can be included in backups, replicas can be added, and so on. The tables are normal Mnesia tables owned by the user of the corresponding `erl_interface` registries.

Exports

`create_table(Tab) -> ok | exit(Reason)`

A wrapper function for `mnesia:create_table/2`, which creates a table (if there is no existing table) with an appropriate set of attributes. The table only resides on the local node and its storage type is the same as the schema table on the local node, that is, `{ram_copies, [node()]}` or `{disc_copies, [node()]}`.

This function is used by `erl_interface` to create the Mnesia table if it does not already exist.

`create_table(Tab, TabDef) -> ok | exit(Reason)`

A wrapper function for `mnesia:create_table/2`, which creates a table (if there is no existing table) with an appropriate set of attributes. The attributes and `TabDef` are forwarded to `mnesia:create_table/2`. For example, if the table is to reside as `disc_only_copies` on all nodes, a call looks as follows:

```
TabDef = [{disc_only_copies, node()|nodes()}],
mnesia_registry:create_table(my_reg, TabDef)
```

See Also

`erl_interface(3)`, `mnesia(3)`