



Runtime_Tools

Copyright © 1999-2021 Ericsson AB. All Rights Reserved.
Runtime_Tools 1.16.2
July 30, 2021

Copyright © 1999-2021 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..

July 30, 2021

1 Runtime Tools User's Guide

Runtime Tools

1.1 LTTng and Erlang/OTP

1.1.1 Introduction

The Linux Trace Toolkit: next generation is an open source system software package for correlated tracing of the Linux kernel, user applications and libraries.

For more information, please visit <http://lttng.org>

1.1.2 Building Erlang/OTP with LTTng support

Configure and build Erlang with LTTng support:

For LTTng to work properly with Erlang/OTP you need the following packages installed:

- LTTng-tools: a command line interface to control tracing sessions.
- LTTng-UST: user space tracing library.

On Ubuntu this can be installed via aptitude:

```
$ sudo aptitude install lttng-tools liblttng-ust-dev
```

See **Installing LTTng** for more information on how to install LTTng on your system.

After LTTng is properly installed on the system Erlang/OTP can be built with LTTng support.

```
$ ./configure --with-dynamic-trace=lttng
$ make
```

1.1.3 Dyntrace Tracepoints

All tracepoints are in the domain of `org_erlang_dyntrace`

All Erlang types are the string equivalent in LTTng.

process_spawn

- `pid` : `string :: Process ID`. Ex. "`<0.131.0>`"
- `parent` : `string :: Process ID`. Ex. "`<0.131.0>`"
- `entry` : `string :: Code Location`. Ex. "`lists:sort/1`"

Available through `erlang:trace/3` with trace flag `procs` and `{tracer, dyntrace, []}` as tracer module.

Example:

```
process_spawn: { cpu_id = 3 }, { pid = "<0.131.0>", parent = "<0.130.0>", entry = "erlang:apply/2" }
```

process_link

- `to` : `string :: Process ID or Port ID`. Ex. "`<0.131.0>`"
- `from` : `string :: Process ID or Port ID`. Ex. "`<0.131.0>`"

1.1 LTTng and Erlang/OTP

- `sbc_carriers` : `integer` :: Number of singleblock carriers in instance. Ex. 1
- `sbc_carriers_size` : `integer` :: Total size of singleblock blocks carriers in instance. Ex. 1343488
- `sbc_blocks` : `integer` :: Number of singleblocks in instance. Ex. 1
- `sbc_blocks_size` : `integer` :: Total size of all singleblock blocks in instance. Ex. 285296

Example:

```
carrier_create: { cpu_id = 2 }, { type = "ets_alloc", instance = 7, size = 2097152, mbc_carriers = 4, mbc_carriers_size = 1343488 }
```

carrier_destroy

- `type` : `string` :: Carrier type. Ex. "ets_alloc"
- `instance` : `integer` :: Allocator instance. Ex. 1
- `size` : `integer` :: Carrier size. Ex. 262144
- `mbc_carriers` : `integer` :: Number of multiblock carriers in instance. Ex. 3
- `mbc_carriers_size` : `integer` :: Total size of multiblock blocks carriers in instance. Ex. 1343488
- `mbc_blocks` : `integer` :: Number of multiblock blocks in instance. Ex. 122
- `mbc_blocks_size` : `integer` :: Total size of all multiblock blocks in instance. Ex. 285296
- `sbc_carriers` : `integer` :: Number of singleblock carriers in instance. Ex. 1
- `sbc_carriers_size` : `integer` :: Total size of singleblock blocks carriers in instance. Ex. 1343488
- `sbc_blocks` : `integer` :: Number of singleblocks in instance. Ex. 1
- `sbc_blocks_size` : `integer` :: Total size of all singleblock blocks in instance. Ex. 285296

Example:

```
carrier_destroy: { cpu_id = 6 }, { type = "ets_alloc", instance = 7, size = 262144, mbc_carriers = 3, mbc_carriers_size = 1343488 }
```

carrier_pool_put

- `type` : `string` :: Carrier type. Ex. "ets_alloc"
- `instance` : `integer` :: Allocator instance. Ex. 1
- `size` : `integer` :: Carrier size. Ex. 262144

Example:

```
carrier_pool_put: { cpu_id = 3 }, { type = "ets_alloc", instance = 5, size = 1048576 }
```

carrier_pool_get

- `type` : `string` :: Carrier type. Ex. "ets_alloc"
- `instance` : `integer` :: Allocator instance. Ex. 1
- `size` : `integer` :: Carrier size. Ex. 262144

Example:

```
carrier_pool_get: { cpu_id = 7 }, { type = "ets_alloc", instance = 4, size = 3208 }
```

1.1.5 Example of process tracing

An example of process tracing of `os_mon` and friends.

Clean start of lttng in a bash shell.

1.2 DTrace and Erlang/OTP

1.2.1 History

The first implementation of DTrace probes for the Erlang virtual machine was presented at the **2008 Erlang User Conference**. That work, based on the Erlang/OTP R12 release, was discontinued due to what appears to be miscommunication with the original developers.

Several users have created Erlang port drivers, linked-in drivers, or NIFs that allow Erlang code to try to activate a probe, e.g. `foo_module:dtrace_probe("message goes here!")`.

1.2.2 Goals

- Annotate as much of the Erlang VM as is practical.
- The initial goal is to trace file I/O operations.
- Support all platforms that implement DTrace: OS X, Solaris, and (I hope) FreeBSD and NetBSD.
- To the extent that it's practical, support SystemTap on Linux via DTrace provider compatibility.
- Allow Erlang code to supply annotations.

1.2.3 Supported platforms

- OS X 10.6.x / Snow Leopard, OS X 10.7.x / Lion and probably newer versions.
- Solaris 10. I have done limited testing on Solaris 11 and OpenIndiana release 151a, and both appear to work.
- FreeBSD 9.0 and 10.0.
- Linux via SystemTap compatibility. Please see `$ERL_TOP/HOWTO/SYSTEMTAP.md` for more details.

Just add the `--with-dynamic-trace=dtrace` option to your command when you run the `configure` script. If you are using `systemtap`, the `configure` option is `--with-dynamic-trace=systemtap`

1.2.4 Status

As of R15B01, the dynamic trace code is included in the OTP source distribution, although it's considered experimental. The main development of the `dtrace` code still happens outside of Ericsson, but there is no need to fetch a patched version of the OTP source to get the basic functionality.

1.2.5 DTrace probe specifications

Probe specifications can be found in `erts/emulator/beam/erlang_dtrace.d`, and a few example scripts can be found under `lib/runtime_tools/examples/`.

1.3 SystemTap and Erlang/OTP

1.3.1 Introduction

SystemTap is DTrace for Linux. In fact Erlang's SystemTap support is built using SystemTap's DTrace compatibility's layer. For an introduction to Erlang DTrace support read `$ERL_TOP/HOWTO/DTRACE.md`.

1.3.2 Requisites

- Linux Kernel with UTRACE support
check for UTRACE support in your current kernel:

```
# grep CONFIG_UTRACE /boot/config-`uname -r`
CONFIG_UTRACE=y
```

Fedora 16 is known to contain UTRACE, for most other Linux distributions a custom build kernel will be required. Check Fedora's SystemTap documentation for additional required packages (e.g. Kernel Debug Symbols)

- SystemTap > 1.6

A the time of writing this, the latest released version of SystemTap is version 1.6. Erlang's DTrace support requires a MACRO that was introduced after that release. So either get a newer release or build SystemTap from git yourself (see: <http://sourceware.org/systemtap/getinvolved.html>)

1.3.3 Building Erlang

Configure and build Erlang with SystemTap support:

```
# ./configure --with-dynamic-trace=systemtap + whatever args you need
# make
```

1.3.4 Testing

SystemTap, unlike DTrace, needs to know what binary it is tracing and has to be able to read that binary before it starts tracing. Your probe script therefor has to reference the correct beam emulator and stap needs to be able to find that binary. The examples are written for "beam", but other versions such as "beam.smp" or "beam.debug.smp" might exist (depending on your configuration). Make sure you either specify the full the path of the binary in the probe or your "beam" binary is in the search path.

All available probes can be listed like this:

```
# stap -L 'process("beam").mark("*")'
```

or:

```
# PATH=/path/to/beam:$PATH stap -L 'process("beam").mark("*")'
```

Probes in the dtrace.so NIF library like this:

```
# PATH=/path/to/dtrace/priv/lib:$PATH stap -L 'process("dtrace.so").mark("*")'
```

1.3.5 Running SystemTap scripts

Adjust the process("beam") reference to your beam version and attach the script to a running "beam" instance:

```
# stap /path/to/probe/script/port1.systemtap -x <pid of beam>
```

2 Reference Manual

Runtime_Tools provides low footprint tracing/debugging tools suitable for inclusion in a production system.

runtime_tools

Application

This chapter describes the Runtime_Tools application in OTP, which provides low footprint tracing/debugging tools suitable for inclusion in a production system.

Configuration

There are currently no configuration parameters available for this application.

SEE ALSO

`application(3)`

dbg

Erlang module

This module implements a text based interface to the `trace/3` and the `trace_pattern/2` BIFs. It makes it possible to trace functions, processes, ports and messages.

To quickly get started on tracing function calls you can use the following code in the Erlang shell:

```
1> dbg:tracer(). %% Start the default trace message receiver
{ok,<0.36.0>}
2> dbg:p(all, c). %% Setup call (c) tracing on all processes
{ok,[{matched,node@nohost,26}]}
3> dbg:tp(lists, seq, x). %% Setup an exception return trace (x) on lists:seq
{ok,[{matched,node@nohost,2},{saved,x}]}
4> lists:seq(1,10).
(<0.34.0>) call lists:seq(1,10)
(<0.34.0>) returned from lists:seq/2 -> [1,2,3,4,5,6,7,8,9,10]
[1,2,3,4,5,6,7,8,9,10]
```

For more examples of how to use `dbg` from the Erlang shell, see the simple example section.

The utilities are also suitable to use in system testing on large systems, where other tools have too much impact on the system performance. Some primitive support for sequential tracing is also included, see the advanced topics section.

Exports

`fun2ms(LiteralFun) -> MatchSpec`

Types:

```
LiteralFun = fun() literal
MatchSpec = term()
```

Pseudo function that by means of a `parse_transform` translates the **literal** `fun()` typed as parameter in the function call to a match specification as described in the `match_spec` manual of ERTS users guide. (With **literal** I mean that the `fun()` needs to textually be written as the parameter of the function, it cannot be held in a variable which in turn is passed to the function).

The parse transform is implemented in the module `ms_transform` and the source **must** include the file `ms_transform.hrl` in `STDLIB` for this pseudo function to work. Failing to include the `hrl` file in the source will result in a runtime error, not a compile time ditto. The include file is easiest included by adding the line `-include_lib("stdlib/include/ms_transform.hrl").` to the source file.

The `fun()` is very restricted, it can take only a single parameter (the parameter list to match), a sole variable or a list. It needs to use the `is_XXX` guard tests and one cannot use language constructs that have no representation in a `match_spec` (like `if`, `case`, `receive` etc). The return value from the `fun` will be the return value of the resulting `match_spec`.

Example:

```
1> dbg:fun2ms(fun([M,N]) when N > 3 -> return_trace() end).
[{['$1','$2'],[{>','$2',3}],[{return_trace}]]]
```

Variables from the environment can be imported, so that this works:

```
2> X=3.
3
3> dbg:fun2ms(fun([M,N]) when N > X -> return_trace() end).
[{'$1','$2'},[{'>','$2',{const,3}}],[{return_trace}]]
```

The imported variables will be replaced by `match_spec` `const` expressions, which is consistent with the static scoping for Erlang `fun()`s. Local or global function calls cannot be in the guard or body of the fun however. Calls to builtin `match_spec` functions of course is allowed:

```
4> dbg:fun2ms(fun([M,N]) when N > X, is_atomm(M) -> return_trace() end).
Error: fun containing local erlang function calls ('is_atomm' called in guard)\
cannot be translated into match_spec
{error,transform_error}
5> dbg:fun2ms(fun([M,N]) when N > X, is_atom(M) -> return_trace() end).
[{'$1','$2'},[{'>','$2',{const,3}},{is_atom,'$1'}],[{return_trace}]]
```

As you can see by the example, the function can be called from the shell too. The `fun()` needs to be literally in the call when used from the shell as well. Other means than the `parse_transform` are used in the shell case, but more or less the same restrictions apply (the exception being records, as they are not handled by the shell).

Warning:

If the `parse_transform` is not applied to a module which calls this pseudo function, the call will fail in runtime (with a `badarg`). The module `dbg` actually exports a function with this name, but it should never really be called except for when using the function in the shell. If the `parse_transform` is properly applied by including the `ms_transform.hrl` header file, compiled code will never call the function, but the function call is replaced by a literal `match_spec`.

More information is provided by the `ms_transform` manual page in `STDLIB`.

`h()` -> ok

Gives a list of items for brief online help.

`h(Item)` -> ok

Types:

Item = **atom()**

Gives a brief help text for functions in the `dbg` module. The available items can be listed with `dbg:h/0`

`p(Item)` -> {ok, MatchDesc} | {error, term()}

Equivalent to `p(Item, [m])`.

`p(Item, Flags)` -> {ok, MatchDesc} | {error, term()}

Types:

MatchDesc = [**MatchNum**]

MatchNum = {**matched**, **node()**, **integer()**} | {**matched**, **node()**, 0, **RPCError**}

RPCError = **term()**

Traces `Item` in accordance to the value specified by `Flags`. The variation of `Item` is listed below:


```
Arity = integer() | '_'
MatchSpec = integer() | Built-inAlias | [] | match_spec()
Built-inAlias = x | c | cx
MatchDesc = [MatchInfo]
MatchInfo = {saved, integer()} | MatchNum
MatchNum = {matched, node(), integer()} | {matched, node(), 0, RPCError}
```

This function enables call trace for one or more functions. All exported functions matching the {Module, Function, Arity} argument will be concerned, but the `match_spec()` may further narrow down the set of function calls generating trace messages.

For a description of the `match_spec()` syntax, please turn to the **User's guide** part of the online documentation for the runtime system (**erts**). The chapter **Match Specifications in Erlang** explains the general match specification "language". The most common generic match specifications used can be found as `Built-inAlias`, see `ltp/0` below for details.

The Module, Function and/or Arity parts of the tuple may be specified as the atom `'_'` which is a "wild-card" matching all modules/functions/arities. Note, if the Module is specified as `'_'`, the Function and Arity parts have to be specified as `'_'` too. The same holds for the Functions relation to the Arity.

All nodes added with `n/1` or `tracer/3` will be affected by this call, and if Module is not `'_'` the module will be loaded on all nodes.

The function returns either an error tuple or a tuple {ok, List}. The List consists of specifications of how many functions that matched, in the same way as the processes and ports are presented in the return value of `p/2`.

There may be a tuple {saved, N} in the return value, if the MatchSpec is other than []. The integer N may then be used in subsequent calls to this function and will stand as an "alias" for the given expression. There are also a couple of built-in aliases for common expressions, see `ltp/0` below for details.

If an error is returned, it can be due to errors in compilation of the match specification. Such errors are presented as a list of tuples {error, string()} where the string is a textual explanation of the compilation error. An example:

```
(x@y)4> dbg:tp({dbg,ltp,0},[[],[],[message, two, arguments], {noexist}])).
{error,
 [{error,"Special form 'message' called with wrong number of
         arguments in {message,two,arguments}."},
  {error,"Function noexist/1 does_not_exist."}]}
```

`tpl(Module,MatchSpec)`

Same as `tpl({Module, '_', '_'}, MatchSpec)`

`tpl(Module,Function,MatchSpec)`

Same as `tpl({Module, Function, '_'}, MatchSpec)`

`tpl(Module, Function, Arity, MatchSpec)`

Same as `tpl({Module, Function, Arity}, MatchSpec)`

`tpl({Module, Function, Arity}, MatchSpec) -> {ok, MatchDesc} | {error, term()}`

This function works as `tp/2`, but enables tracing for local calls (and local functions) as well as for global calls (and functions).

```
tpe(Event, MatchSpec) -> {ok, MatchDesc} | {error, term()}
```

Types:

```
Event = send | 'receive'
MatchSpec = integer() | Built-inAlias | [] | match_spec()
Built-inAlias = x | c | cx
MatchDesc = [MatchInfo]
MatchInfo = {saved, integer()} | MatchNum
MatchNum = {matched, node(), 1} | {matched, node(), 0, RPCError}
```

This function associates a match specification with trace event `send` or `'receive'`. By default all executed `send` and `'receive'` events are traced if enabled for a process. A match specification can be used to filter traced events based on sender, receiver and/or message content.

For a description of the `match_spec()` syntax, please turn to the **User's guide** part of the online documentation for the runtime system (**erts**). The chapter **Match Specifications in Erlang** explains the general match specification "language".

For `send`, the matching is done on the list `[Receiver, Msg]`. `Receiver` is the process or port identity of the receiver and `Msg` is the message term. The pid of the sending process can be accessed with the guard function `self/0`.

For `'receive'`, the matching is done on the list `[Node, Sender, Msg]`. `Node` is the node name of the sender. `Sender` is the process or port identity of the sender, or the atom `undefined` if the sender is not known (which may be the case for remote senders). `Msg` is the message term. The pid of the receiving process can be accessed with the guard function `self/0`.

All nodes added with `n/1` or `tracer/3` will be affected by this call.

The return value is the same as for `tp/2`. The number of matched events are never larger than 1 as `tpe/2` does not accept any form of wildcards for argument `Event`.

```
ctp()
```

Same as `ctp({'_', '_', '_'})`

```
ctp(Module)
```

Same as `ctp({Module, '_', '_'})`

```
ctp(Module, Function)
```

Same as `ctp({Module, Function, '_'})`

```
ctp(Module, Function, Arity)
```

Same as `ctp({Module, Function, Arity})`

```
ctp({Module, Function, Arity}) -> {ok, MatchDesc} | {error, term()}
```

Types:

```
Module = atom() | '_'
Function = atom() | '_'
Arity = integer() | '_'
MatchDesc = [MatchNum]
MatchNum = {matched, node(), integer()} | {matched, node(), 0, RPCError}
```

This function disables call tracing on the specified functions. The semantics of the parameter is the same as for the corresponding function specification in `tp/2` or `tpl/2`. Both local and global call trace is disabled.

The return value reflects how many functions that matched, and is constructed as described in `tp/2`. No tuple `{saved, N}` is however ever returned (for obvious reasons).

`ctpl()`

Same as `ctpl({'_', '_', '_'})`

`ctpl(Module)`

Same as `ctpl({Module, '_', '_'})`

`ctpl(Module, Function)`

Same as `ctpl({Module, Function, '_'})`

`ctpl(Module, Function, Arity)`

Same as `ctpl({Module, Function, Arity})`

`ctpl({Module, Function, Arity}) -> {ok, MatchDesc} | {error, term()}`

This function works as `ctp/1`, but only disables tracing set up with `tpl/2` (not with `tp/2`).

`ctpg()`

Same as `ctpg({'_', '_', '_'})`

`ctpg(Module)`

Same as `ctpg({Module, '_', '_'})`

`ctpg(Module, Function)`

Same as `ctpg({Module, Function, '_'})`

`ctpg(Module, Function, Arity)`

Same as `ctpg({Module, Function, Arity})`

`ctpg({Module, Function, Arity}) -> {ok, MatchDesc} | {error, term()}`

This function works as `ctp/1`, but only disables tracing set up with `tp/2` (not with `tpl/2`).

`ctpe(Event) -> {ok, MatchDesc} | {error, term()}`

Types:

`Event = send | 'receive'`

`MatchDesc = [MatchNum]`

`MatchNum = {matched, node(), 1} | {matched, node(), 0, RPCError}`

This function clears match specifications for the specified trace event (`send` or `'receive'`). It will revert back to the default behavior of tracing all triggered events.

The return value follow the same style as for `ctp/1`.


```

Type = port | process | module
Data = PortGenerator | HandlerSpec | ModuleSpec
PortGenerator = fun() (no arguments)
Error = term()
HandlerSpec = {HandlerFun, InitialData}
HandlerFun = fun() (two arguments)
ModuleSpec = fun() (no arguments) | {TracerModule, TracerState}
TracerModule = atom()
InitialData = TracerState = term()

```

This function starts a tracer server with additional parameters on the local node. The first parameter, the `Type`, indicates if trace messages should be handled by a receiving process (`process`), by a tracer port (`port`) or by a tracer module (`module`). For a description about tracer ports see `trace_port/2` and for a tracer modules see `erl_tracer`.

If `Type` is `process`, a message handler function can be specified (`HandlerSpec`). The handler function, which should be a `fun` taking two arguments, will be called for each trace message, with the first argument containing the message as it is and the second argument containing the return value from the last invocation of the `fun`. The initial value of the second parameter is specified in the `InitialData` part of the `HandlerSpec`. The `HandlerFun` may choose any appropriate action to take when invoked, and can save a state for the next invocation by returning it.

If `Type` is `port`, then the second parameter should be a **fun** which takes no arguments and returns a newly opened trace port when called. Such a **fun** is preferably generated by calling `trace_port/2`.

if `Type` is `module`, then the second parameter should be either a tuple describing the `erl_tracer` module to be used for tracing and the state to be used for that tracer module or a `fun` returning the same tuple.

If an error is returned, it can either be due to a tracer server already running (`{error, already_started}`) or due to the `HandlerFun` throwing an exception.

To start a similar tracer on a remote node, use `tracer/3`.

```
tracer(Nodename, Type, Data) -> {ok, Nodename} | {error, Reason}
```

Types:

```
Nodename = atom()
```

This function is equivalent to `tracer/2`, but acts on the given node. A tracer is started on the node (`Nodename`) and the node is added to the list of traced nodes.

Note:

This function is not equivalent to `n/1`. While `n/1` starts a process tracer which redirects all trace information to a process tracer on the local node (i.e. the trace control node), `tracer/3` starts a tracer of any type which is independent of the tracer on the trace control node.

For details, see `tracer/2`.

```
trace_port(Type, Parameters) -> fun()
```

Types:

```

Type = ip | file
Parameters = Filename | WrapFilesSpec | IPPortSpec
Filename = string() | [string()] | atom()

```



```
ant@stack> dbg:tracer(port, dbg:trace_port(ip,4711)).
<0.17.0>
ant@stack> dbg:p(self(), send).
{ok,1}
```

All trace messages are now sent to the trace port driver, which in turn listens for connections on the TCP/IP port 4711. If we want to see the messages on another node, preferably on another host, we do like this:

```
-> dbg:trace_client(ip, {"stack", 4711}).
<0.42.0>
```

If we now send a message from the shell on the node `ant@stack`, where all sends from the shell are traced:

```
ant@stack> self() ! hello.
hello
```

The following will appear at the console on the node that started the trace client:

```
(<0.23.0>) <0.23.0> ! hello
(<0.23.0>) <0.22.0> ! {shell_rep,<0.23.0>,{value,hello,[],[]}}
```

The last line is generated due to internal message passing in the Erlang shell. The process id's will vary.

`trace_client(Type, Parameters, HandlerSpec) -> pid()`

Types:

```
Type = ip | file | follow_file
Parameters = Filename | WrapFilesSpec | IPClientPortSpec
Filename = string() | [string()] | atom()
WrapFilesSpec = see trace_port/2
Suffix = string()
IPClientPortSpec = PortNumber | {Hostname, PortNumber}
PortNumber = integer()
Hostname = string()
HandlerSpec = {HandlerFun, InitialData}
HandlerFun = fun() (two arguments)
InitialData = term()
```

This function works exactly as `trace_client/2`, but allows you to write your own handler function. The handler function works mostly as the one described in `tracer/2`, but will also have to be prepared to handle trace messages of the form `{drop, N}`, where `N` is the number of dropped messages. This pseudo trace message will only occur if the `ip` trace driver is used.

For trace type `file`, the pseudo trace message `end_of_trace` will appear at the end of the trace. The return value from the handler function is in this case ignored.

`stop_trace_client(Pid) -> ok`

Types:

```
Pid = pid()
```

This function shuts down a previously started trace client. The `Pid` argument is the process id returned from the `trace_client/2` or `trace_client/3` call.

`get_tracer()`

Equivalent to `get_tracer(node())`.

`get_tracer(Nodename) -> {ok, Tracer}`

Types:

```
Nodename = atom()
Tracer = port() | pid() | {module(), term()}
```

Returns the process, port or tracer module to which all trace messages are sent.

`stop() -> ok`

Stops the dbg server and clears all trace flags for all processes and all local trace patterns for all functions. Also shuts down all trace clients and closes all trace ports.

Note that no global trace patterns are affected by this function.

`stop_clear() -> ok`

Same as `stop/0`, but also clears all trace patterns on global functions calls.

Simple examples - tracing from the shell

The simplest way of tracing from the Erlang shell is to use `dbg:c/3` or `dbg:c/4`, e.g. tracing the function `dbg:get_tracer/0`:

```
(tiger@durin)84> dbg:c(dbg,get_tracer,[]).
(<0.154.0>) <0.152.0> ! {<0.154.0>,{get_tracer,tiger@durin}}
(<0.154.0>) out {dbg,req,1}
(<0.154.0>) << {dbg,{ok,<0.153.0>}}
(<0.154.0>) in {dbg,req,1}
(<0.154.0>) << timeout
{ok,<0.153.0>}
(tiger@durin)85>
```

Another way of tracing from the shell is to explicitly start a **tracer** and then set the **trace flags** of your choice on the processes you want to trace, e.g. trace messages and process events:

```
(tiger@durin)66> Pid = spawn(fun() -> receive {From,Msg} -> From ! Msg end end).
<0.126.0>
(tiger@durin)67> dbg:tracer().
{ok,<0.128.0>}
(tiger@durin)68> dbg:p(Pid,[m,procs]).
{ok,[{matched,tiger@durin,1}]}
(tiger@durin)69> Pid ! {self(),hello}.
(<0.126.0>) << {<0.116.0>,hello}
{<0.116.0>,hello}
(<0.126.0>) << timeout
(<0.126.0>) <0.116.0> ! hello
(<0.126.0>) exit normal
(tiger@durin)70> flush().
Shell got hello
ok
(tiger@durin)71>
```

If you set the `call` trace flag, you also have to set a **trace pattern** for the functions you want to trace:

```
(tiger@durin)77> dbg:tracer().
{ok,<0.142.0>}
(tiger@durin)78> dbg:p(all,call).
{ok,[{matched,tiger@durin,3}]}
(tiger@durin)79> dbg:tp(dbg,get_tracer,0,[]).
{ok,[{matched,tiger@durin,1}]}
(tiger@durin)80> dbg:get_tracer().
(<0.116.0>) call dbg:get_tracer()
{ok,<0.143.0>}
(tiger@durin)81> dbg:tp(dbg,get_tracer,0,[{'_',[],[{return_trace}]}]).
{ok,[{matched,tiger@durin,1},{saved,1}]}
(tiger@durin)82> dbg:get_tracer().
(<0.116.0>) call dbg:get_tracer()
(<0.116.0>) returned from dbg:get_tracer/0 -> {ok,<0.143.0>}
{ok,<0.143.0>}
(tiger@durin)83>
```

Advanced topics - combining with seq_trace

The `dbg` module is primarily targeted towards tracing through the `erlang:trace/3` function. It is sometimes desired to trace messages in a more delicate way, which can be done with the help of the `seq_trace` module.

`seq_trace` implements sequential tracing (known in the AXE10 world, and sometimes called "forlopp tracing"). `dbg` can interpret messages generated from `seq_trace` and the same tracer function for both types of tracing can be used. The `seq_trace` messages can even be sent to a trace port for further analysis.

As a match specification can turn on sequential tracing, the combination of `dbg` and `seq_trace` can be quite powerful. This brief example shows a session where sequential tracing is used:

```

1> dbg:tracer().
{ok,<0.30.0>}
2> {ok, Tracer} = dbg:get_tracer().
{ok,<0.31.0>}
3> seq_trace:set_system_tracer(Tracer).
false
4> dbg:tp(dbg, get_tracer, 0, [[[],[]],[{set_seq_token, send, true}]]).
{ok,[{matched,nonode@nohost,1},{saved,1}]}
5> dbg:p(all,call).
{ok,[{matched,nonode@nohost,22}]}
6> dbg:get_tracer(), seq_trace:set_token([]).
(<0.25.0>) call dbg:get_tracer()
SeqTrace [0]: (<0.25.0>) <0.30.0> ! {<0.25.0>,get_tracer} [Serial: {2,4}]
SeqTrace [0]: (<0.30.0>) <0.25.0> ! {dbg,{ok,<0.31.0>}} [Serial: {4,5}]
{1,0,5,<0.30.0>,4}

```

This session sets the `system_tracer` to the same process as the ordinary tracer process (i. e. `<0.31.0>`) and sets the trace pattern for the function `dbg:get_tracer` to one that has the action of setting a sequential token. When the function is called by a traced process (all processes are traced in this case), the process gets "contaminated" by the token and `seq_trace` messages are sent both for the server request and the response. The `seq_trace:set_token([])` after the call clears the `seq_trace` token, why no messages are sent when the answer propagates via the shell to the console port. The output would otherwise have been more noisy.

Note of caution

When tracing function calls on a group leader process (an IO process), there is risk of causing a deadlock. This will happen if a group leader process generates a trace message and the tracer process, by calling the trace handler function, sends an IO request to the same group leader. The problem can only occur if the trace handler prints to `tty` using an `io` function such as `format/2`. Note that when `dbg:p(all,call)` is called, IO processes are also traced. Here's an example:

```

%% Using a default line editing shell
1> dbg:tracer(process, {fun(Msg,_) -> io:format("~p~n", [Msg]), 0 end, 0}).
{ok,<0.37.0>}
2> dbg:p(all, [call]).
{ok,[{matched,nonode@nohost,25}]}
3> dbg:tp(mymod,[{'_',[],[]}]).
{ok,[{matched,nonode@nohost,0},{saved,1}]}
4> mymod: % TAB pressed here
%% -- Deadlock --

```

Here's another example:

```

%% Using a shell without line editing (oldshell)
1> dbg:tracer(process).
{ok,<0.31.0>}
2> dbg:p(all, [call]).
{ok,[{matched,nonode@nohost,25}]}
3> dbg:tp(lists,[{'_',[],[]}]).
{ok,[{matched,nonode@nohost,0},{saved,1}]}
% -- Deadlock --

```

The reason we get a deadlock in the first example is because when `TAB` is pressed to expand the function name, the group leader (which handles character input) calls `mymod:module_info()`. This generates a trace message which, in turn, causes the tracer process to send an IO request to the group leader (by calling `io:format/2`). We end up in a deadlock.

In the second example we use the default trace handler function. This handler prints to `tty` by sending IO requests to the user process. When Erlang is started in oldshell mode, the shell process will have `user` as its group leader and so will the tracer process in this example. Since `user` calls functions in `lists` we end up in a deadlock as soon as the first IO request is sent.

Here are a few suggestions for how to avoid deadlock:

- Don't trace the group leader of the tracer process. If tracing has been switched on for all processes, call `dbg:p(TracerGLPid,clear)` to stop tracing the group leader (`TracerGLPid`). `process_info(TracerPid,group_leader)` tells you which process this is (`TracerPid` is returned from `dbg:get_tracer/0`).
- Don't trace the user process if using the default trace handler function.
- In your own trace handler function, call `erlang:display/1` instead of an `io` function or, if `user` is not used as group leader, print to `user` instead of the default group leader. Example:
`io:format(user,Str,Args).`

dyntrace

Erlang module

This module implements interfaces to dynamic tracing, should such be compiled into the virtual machine. For a standard and/or commercial build, no dynamic tracing is available, in which case none of the functions in this module is usable or give any effect.

Should dynamic tracing be enabled in the current build, either by configuring with `./configure --with-dynamic-trace=dtrace` or with `./configure --with-dynamic-trace=systemtap`, the module can be used for two things:

- Trigger the user-probe `user_trace_i4s4` in the NIF library `dyntrace.so` by calling `dyntrace:p/{1,2,3,4,5,6,7,8}`.
- Set a user specified tag that will be present in the trace messages of both the `efile_drv` and the user-probe mentioned above.

Both building with dynamic trace probes and using them is experimental and unsupported by Erlang/OTP. It is included as an option for the developer to trace and debug performance issues in their systems.

The original implementation is mostly done by Scott Lystiger Fritchie as an Open Source Contribution and it should be viewed as such even though the source for dynamic tracing as well as this module is included in the main distribution. However, the ability to use dynamic tracing of the virtual machine is a very valuable contribution which OTP has every intention to maintain as a tool for the developer.

How to write `d` programs or `systemtap` scripts can be learned from books and from a lot of pages on the Internet. This manual page does not include any documentation about using the dynamic trace tools of respective platform. The `examples` directory of the `runtime_tools` application however contains comprehensive examples of both `d` and `systemtap` programs that will help you get started. Another source of information is the `dtrace` and `systemtap` chapters in the Runtime Tools Users' Guide.

Exports

`available() -> boolean()`

This function uses the NIF library to determine if dynamic tracing is available. Usually calling `erlang:system_info/1` is a better indicator of the availability of dynamic tracing.

The function will throw an exception if the `dyntrace` NIF library could not be loaded by the `on_load` function of this module.

`p() -> true | false | error | badarg`

Calling this function will trigger the "user" trace probe `user_trace_i4s4` in the `dyntrace` NIF module, sending a trace message only containing the user tag and zeroes/empty strings in all other fields.

`p(integer() | string()) -> true | false | error | badarg`

Calling this function will trigger the "user" trace probe `user_trace_i4s4` in the `dyntrace` NIF module, sending a trace message containing the user tag and the integer or string parameter in the first integer/string field.

```
p(integer() | string(), integer() | string()) -> true | false | error | badarg
```

Calling this function will trigger the "user" trace probe user_trace_i4s4 in the dyntrace NIF module, sending a trace message containing the user tag and the integer() or string() parameters as the first fields of respective type. integer() parameters should be put before any string() parameters. I.e. p(1, "Hello") is ok, as is p(1,1) and p("Hello", "Again"), but not p("Hello",1).

```
p(integer() | string(), integer() | string(), integer() | string()) -> true | false | error | badarg
```

Calling this function will trigger the "user" trace probe user_trace_i4s4 in the dyntrace NIF module, sending a trace message containing the user tag and the integer() or string() parameters as the first fields of respective type. integer() parameters should be put before any string() parameters, as in p/2.

```
p(integer() | string(), integer() | string(), integer() | string(), integer() | string()) -> true | false | error | badarg
```

Calling this function will trigger the "user" trace probe user_trace_i4s4 in the dyntrace NIF module, sending a trace message containing the user tag and the integer() or string() parameters as the first fields of respective type. integer() parameters should be put before any string() parameters, as in p/2.

```
p(integer(), integer() | string(), integer() | string(), integer() | string(), string(), string()) -> true | false | error | badarg
```

Calling this function will trigger the "user" trace probe user_trace_i4s4 in the dyntrace NIF module, sending a trace message containing the user tag and the integer() or string() parameters as the first fields of respective type. integer() parameters should be put before any string() parameters, as in p/2.

There can be no more than four parameters of any type (integer() or string()), so the first parameter has to be an integer() and the last a string().

```
p(integer(), integer(), integer() | string(), integer() | string(), string(), string(), string()) -> true | false | error | badarg
```

Calling this function will trigger the "user" trace probe user_trace_i4s4 in the dyntrace NIF module, sending a trace message containing the user tag and the integer() or string() parameters as the first fields of respective type. integer() parameters should be put before any string() parameters, as in p/2.

There can be no more than four parameters of any type (integer() or string()), so the first two parameters has to be integer()'s and the last two string()'s.

```
p(integer(), integer(), integer(), integer() | string(), string(), string(), string(), string()) -> true | false | error | badarg
```

Calling this function will trigger the "user" trace probe user_trace_i4s4 in the dyntrace NIF module, sending a trace message containing the user tag and the integer() or string() parameters as the first fields of respective type. integer() parameters should be put before any string() parameters, as in p/2.

There can be no more than four parameters of any type (integer() or string()), so the first three parameters has to be integer()'s and the last three string()'s.

In this example, any user tag set in the calling process will be spread to the I/O-server when the `io:format` call is done.

```
restore_tag(TagData) -> true
```

Types:

TagData = opaque data returned by spread_tag/1

Restores the previous state of user tags and their spreading as it was before a call to `spread_tag/1`. Note that the restoring is not limited to the same process, one can utilize this to turn off spreading in one process and restore it in a newly created, the one that actually is going to send messages:

```
f() ->
  TagData=dyntrace:spread_tag(false),
  spawn(fun() ->
    dyntrace:restore_tag(TagData),
    do_something()
  end),
  do_something_else(),
  dyntrace:restore_tag(TagData).
```

Correctly handling user tags and their spreading might take some effort, as Erlang programs tend to send and receive messages so that sometimes the user tag gets lost due to various things, like double receives or communication with a port (ports do not handle user tags, in the same way as they do not handle regular sequential trace tokens).

erts_alloc_config

Erlang module

Note:

`erts_alloc_config` is currently an experimental tool and might be subject to backward incompatible changes.

`erts_alloc(3)` is an Erlang Run-Time System internal memory allocator library. `erts_alloc_config` is intended to be used to aid creation of an `erts_alloc(3)` configuration that is suitable for a limited number of runtime scenarios. The configuration that `erts_alloc_config` produce is intended as a suggestion, and may need to be adjusted manually.

The configuration is created based on information about a number of runtime scenarios. It is obviously impossible to foresee every runtime scenario that can occur. The important scenarios are those that cause maximum or minimum load on specific memory allocators. Load in this context is total size of memory blocks allocated.

The current implementation of `erts_alloc_config` concentrate on configuration of multi-block carriers. Information gathered when a runtime scenario is saved is mainly current and maximum use of multi-block carriers. If a parameter that change the use of multi-block carriers is changed, a previously generated configuration is invalid and `erts_alloc_config` needs to be run again. It is mainly the single block carrier threshold that effects the use of multi-block carriers, but other single-block carrier parameters might as well. If another value of a single block carrier parameter than the default is desired, use the desired value when running `erts_alloc_config`.

A configuration is created in the following way:

- Pass the `+Mea config` command-line flag to the Erlang runtime system you are going to use for creation of the allocator configuration. It will disable features that prevent `erts_alloc_config` from doing its job. Note, you should **not** use this flag when using the created configuration. Also note that it is important that you use the same amount of schedulers when creating the configuration as you are going the use on the system using the configuration.
- Run your applications with different scenarios (the more the better) and save information about each scenario by calling `save_scenario/0`. It may be hard to know when the applications are at an (for `erts_alloc_config`) important runtime scenario. A good approach may therefore be to call `save_scenario/0` repeatedly, e.g. once every tenth second. Note that it is important that your applications reach the runtime scenarios that are important for `erts_alloc_config` when you are saving scenarios; otherwise, the configuration may perform bad.
- When you have covered all scenarios, call `make_config/1` in order to create a configuration. The configuration is written to a file that you have chosen. This configuration file can later be read by an Erlang runtime-system at startup. Pass the command line argument `-args_file FileName` to the `erl(1)` command.
- The configuration produced by `erts_alloc_config` may need to be manually adjusted as already stated. Do not modify the file produced by `erts_alloc_config`; instead, put your modifications in another file and load this file after the file produced by `erts_alloc_config`. That is, put the `-args_file FileName` argument that reads your modification file later on the command-line than the `-args_file FileName` argument that reads the configuration file produced by `erts_alloc_config`. If a memory allocation parameter appear multiple times, the last version of will be used, i.e., you can override parameters in the configuration file produced by `erts_alloc_config`. Doing it this way simplifies things when you want to rerun `erts_alloc_config`.


```
    id := msacc_id(),
    counters := msacc_data_counters()}
msacc_data_counters() = #{msacc_state() => integer() >= 0}
```

A map containing the different microstate accounting states and the number of microseconds spent in it.

```
msacc_stats() = [msacc_stats_thread()]
msacc_stats_thread() =
    #{'$type' := msacc_stats,
     type := msacc_type(),
     id := msacc_id(),
     system := float(),
     counters := msacc_stats_counters()}
```

A map containing information about a specific thread. The percentages in the map can be either run-time or real-time depending on if `runtime` or `realtime` was requested from `stats/2`. `system` is the percentage of total system time for this specific thread.

```
msacc_stats_counters() =
    #{msacc_state() => #{thread := float(), system := float()}}
```

A map containing the different microstate accounting states. Each value in the map contains another map with the percentage of time that this thread has spent in the specific state. Both the percentage of `system` time and the time for that specific thread is part of the map.

```
msacc_type() =
    aux | async | dirty_cpu_scheduler | dirty_io_scheduler |
    poll | scheduler
msacc_id() = integer() >= 0
msacc_state() =
    alloc | aux | bif | busy_wait | check_io | emulator | ets |
    gc | gc_fullsweep | nif | other | port | send | sleep | timers
```

The different states that a thread can be in. See `erlang:statistics(microstate_accounting)` for details.

```
msacc_print_options() = #{system => boolean()}
```

The different options that can be given to `print/2`.

Exports

```
available() -> boolean()
```

This function checks whether microstate accounting is available or not.

```
start() -> boolean()
```

Start microstate accounting. Returns whether it was previously enabled or disabled.

```
start(Time) -> true
```

Types:

```
    Time = timeout()
```

Resets all counters and then starts microstate accounting for the given milliseconds.


```
stats(Analysis, Stats) -> integer() >= 0
```

Types:

```
Analysis = system_realtime | system_runtime
```

```
Stats = msacc_data()
```

Returns the system time for the given microstate statistics values. System time is the accumulated time of all threads.

`realtime`

Returns all time recorded for all threads.

`runtime`

Returns all time spent doing work for all threads, i.e. all time not spent in the `sleep` state.

```
stats(Analysis, Stats) -> msacc_stats()
```

Types:

```
Analysis = realtime | runtime
```

```
Stats = msacc_data()
```

Returns fractions of real-time or run-time spent in the various threads from the given microstate statistics values.

```
stats(Analysis, StatsOrData) -> msacc_data() | msacc_stats()
```

Types:

```
Analysis = type
```

```
StatsOrData = msacc_data() | msacc_stats()
```

Returns a list of microstate statistics values where the values for all threads of the same type has been merged.

```
to_file(Filename) -> ok | {error, file:posix()}
```

Types:

```
Filename = file:name_all()
```

Dumps the current microstate statistics counters to a file that can be parsed with `file:consult/1`.

```
from_file(Filename) -> msacc_data()
```

Types:

```
Filename = file:name_all()
```

Read a file dump produced by `to_file(Filename)`.

scheduler

Erlang module

This module contains utility functions for easier measurement and calculation of scheduler utilization, otherwise obtained from calling the more primitive `statistics(scheduler_wall_time)`.

The simplest usage is to call `scheduler:utilization(Seconds)`.

Data Types

```

sched_sample()
sched_type() = normal | cpu | io
sched_id() = integer()
sched_util_result() =
    [{sched_type(), sched_id(), float(), string()} |
     {total, float(), string()} |
     {weighted, float(), string()}]

```

A list of tuples containing results for individual schedulers as well as aggregated averages. `Util` is the scheduler utilization as a floating point value between 0.0 and 1.0. `Percent` is the same utilization as a more human readable string expressed in percent.

```

{normal, SchedulerId, Util, Percent}
    Scheduler utilization of a normal scheduler with number SchedulerId. Schedulers that are not online will
    also be included. Online schedulers have the lowest SchedulerId.
{cpu, SchedulerId, Util, Percent}
    Scheduler utilization of a dirty-cpu scheduler with number SchedulerId.
{io, SchedulerId, Util, Percent}
    Scheduler utilization of a dirty-io scheduler with number SchedulerId. This tuple will only exist if both
    samples were taken with sample_all/0.
{total, Util, Percent}
    Total utilization of all normal and dirty-cpu schedulers.
{weighted, Util, Percent}
    Total utilization of all normal and dirty-cpu schedulers, weighted against maximum amount of available CPU
    time.

```

Exports

`sample()` -> `sched_sample()`

Return a scheduler utilization sample for normal and dirty-cpu schedulers.

`sample_all()` -> `sched_sample()`

Return a scheduler utilization sample for all schedulers, including dirty-io schedulers.

`utilization(Seconds)` -> `sched_util_result()`

Types:

`Seconds = integer() >= 1`

Measure utilization for normal and dirty-cpu schedulers during `Seconds` seconds, and then return the result.

```
utilization(Sample) -> sched_util_result()
```

Types:

```
Sample = sched_sample()
```

Calculate scheduler utilizations for the time interval from when `Sample` was taken and "now". The same as calling `scheduler:utilization(Sample, scheduler:sample_all())`.

Note:

Scheduler utilization is measured as an average value over a time interval, calculated as the difference between two samples. To get good useful utilization values at least a couple of seconds should have passed between the two samples. For this reason, you should not do

```
scheduler:utilization(scheduler:sample()). % DO NOT DO THIS!
```

The above example takes two samples in rapid succession and calculates the scheduler utilization between them. The resulting values will probably be more misleading than informative.

Instead use `scheduler:utilization(Seconds)` or let some time pass between `Sample=scheduler:sample()` and `scheduler:utilization(Sample)`.

```
utilization(Sample1, Sample2) -> sched_util_result()
```

Types:

```
Sample1 = Sample2 = sched_sample()
```

Calculates scheduler utilizations for the time interval between the two samples obtained from calling `sample/0` or `sample_all/0`.

`to_file(FileName) -> ok | {error, Reason}`

Types:

`FileName = file:name_all()`

`Reason = file:posix() | badarg | terminated | system_limit`

Writes miscellaneous system information to file. This information will typically be requested by the Erlang/OTP team at Ericsson AB when reporting an issue.