



Runtime_Tools

Copyright © 1999-2024 Ericsson AB. All Rights Reserved.
Runtime_Tools 1.19
December 5, 2024

Copyright © 1999-2024 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..

December 5, 2024

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

At 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 therefore 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:


```
(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:

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).`

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

Warning:

This (experimental) tool no longer produces good configurations and cannot be fixed in a reasonably backwards compatible manner. It has therefore been scheduled for removal in OTP 26.0.

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.


```
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)`.

Note:

This function is **not recommended** as it's so easy to get invalid results without noticing. In particular do not do this:

```
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/2` and call `get_sample/0` to get samples with some time in between.

```
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 `get_sample/0` or `get_sample_all/0`.

This function itself, does not need `scheduler_wall_time` to be enabled. However, for a correct result, `scheduler_wall_time` must have been enabled during the entire interval between the two samples.

system_information

Erlang module

Exports

`sanity_check() -> ok | {failed, Failures}`

Types:

```
Application = atom()
ApplicationVersion = string()
MissingRuntimeDependencies =
    {missing_runtime_dependencies, ApplicationVersion,
     [ApplicationVersion]}
InvalidApplicationVersion =
    {invalid_application_version, ApplicationVersion}
InvalidAppFile = {invalid_app_file, Application}
Failure =
    MissingRuntimeDependencies | InvalidApplicationVersion |
    InvalidAppFile
Failures = [Failure]
```

Performs a sanity check on the system. If no issues were found, `ok` is returned. If issues were found, `{failed, Failures}` is returned. All failures found will be part of the `Failures` list. Currently defined `Failure` elements in the `Failures` list:

`InvalidAppFile`

An application has an invalid `.app` file. The second element identifies the application which has the invalid `.app` file.

`InvalidApplicationVersion`

An application has an invalid application version. The second element identifies the application version that is invalid.

`MissingRuntimeDependencies`

An application is missing runtime dependencies. The second element identifies the application (with version) that has missing dependencies. The third element contains the missing dependencies.

Note that this check use application versions that are loaded, or will be loaded when used. You might have application versions that satisfies all dependencies installed in the system, but if those are not loaded this check will fail. The system will of course also fail when used like this. This may happen when you have multiple branched versions of the same application installed in the system, but you do not use a boot script identifying the correct application version.

Currently the sanity check is limited to verifying runtime dependencies found in the `.app` files of all applications. More checks will be introduced in the future. This implies that the return type **will** change in the future.

Note:

An `ok` return value only means that `sanity_check/0` did not find any issues, **not** that no issues exist.

```
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.