

Erlang Run-Time System Application (ERTS)

version 5.6

Typeset in L^AT_EX from SGML source using the DocBuilder-0.9.8.4 Document System.

Contents

1	ERTS User's Guide	1
1.1	Match specifications in Erlang	1
1.1.1	Grammar	1
1.1.2	Function descriptions	2
1.1.3	Variables and literals	4
1.1.4	Execution of the match	5
1.1.5	Differences between match specifications in ETS and tracing	5
1.1.6	Examples	6
1.2	How to interpret the Erlang crash dumps	7
1.2.1	General information	8
1.2.2	Memory information	9
1.2.3	Internal table information	9
1.2.4	Allocated areas	9
1.2.5	Allocator	9
1.2.6	Process information	9
1.2.7	Port information	10
1.2.8	ETS tables	11
1.2.9	Timers	11
1.2.10	Distribution information	11
1.2.11	Loaded module information	12
1.2.12	Fun information	12
1.2.13	Process Data	12
1.2.14	Atoms	13
1.2.15	Disclaimer	13
1.3	How to implement an alternative carrier for the Erlang distribution	13
1.3.1	Introduction	13
1.3.2	The driver	14
1.3.3	Putting it all together	28
1.4	The Abstract Format	29
1.4.1	Module declarations and forms	29

1.4.2	Atomic literals	30
1.4.3	Patterns	30
1.4.4	Expressions	31
1.4.5	Clauses	33
1.4.6	Guards	34
1.4.7	The abstract format after preprocessing	34
1.5	tty - A command line interface	35
1.5.1	Normal Mode	35
1.5.2	Shell Break Mode	36
1.6	How to implement a driver	36
1.6.1	Introduction	36
1.6.2	Sample driver	36
1.6.3	Compiling and linking the sample driver	41
1.6.4	Calling a driver as a port in Erlang	41
1.6.5	Sample asynchronous driver	42
1.6.6	An asynchronous driver using driver_async	45
1.7	Inet configuration	48
1.7.1	Introduction	48
1.7.2	Configuration Data	49
1.7.3	User Configuration Example	50
1.8	External Term Format	51
1.8.1	Introduction	51
1.8.2	SMALL_INTEGER_EXT	52
1.8.3	INTEGER_EXT	52
1.8.4	FLOAT_EXT	52
1.8.5	ATOM_EXT	53
1.8.6	REFERENCE_EXT	53
1.8.7	PORT_EXT	53
1.8.8	PID_EXT	54
1.8.9	SMALL_TUPLE_EXT	54
1.8.10	LARGE_TUPLE_EXT	54
1.8.11	NIL_EXT	54
1.8.12	STRING_EXT	55
1.8.13	LIST_EXT	55
1.8.14	BINARY_EXT	55
1.8.15	SMALL_BIG_EXT	55
1.8.16	LARGE_BIG_EXT	56
1.8.17	NEW_CACHE	56
1.8.18	CACHED_ATOM	56
1.8.19	NEW_REFERENCE_EXT	57

1.8.20	FUN_EXT	57
1.8.21	NEW_FUN_EXT	57
1.8.22	EXPORT_EXT	58
1.8.23	BIT_BINARY_EXT	58
1.8.24	NEW_FLOAT_EXT	59
1.9	Distribution Protocol	59
1.9.1	EPMD Protocol	59
1.9.2	Handshake	65
1.9.3	Protocol between connected nodes	65
1.9.4	New Ctrlmessages for distrvsn = 1 (OTP R4)	66
1.9.5	New Ctrlmessages for distrvsn = 2	66
1.9.6	New Ctrlmessages for distrvsn = 3 (OTP R5C)	66
1.9.7	New Ctrlmessages for distrvsn = 4 (OTP R6)	66
2	ERTS Reference Manual	69
2.1	epmd	78
2.2	erl	79
2.3	erlc	86
2.4	erlsrv	89
2.5	escript	94
2.6	run_erl	96
2.7	start	98
2.8	start_erl	99
2.9	werl	101
2.10	erl_set_memory_block	102
2.11	erts_alloc	104
2.12	driver_entry	110
2.13	erl_driver	116
	List of Figures	147
	List of Tables	149

Chapter 1

ERTS User's Guide

The Erlang Runtime System Application *ERTS*.

1.1 Match specifications in Erlang

A “match specification” (*match_spec*) is an Erlang term describing a small “program” that will try to match something (either the parameters to a function as used in the `erlang:trace_pattern/2` BIF, or the objects in an ETS table.). The *match_spec* in many ways works like a small function in Erlang, but is interpreted/compiled by the Erlang runtime system to something much more efficient than calling an Erlang function. The *match_spec* is also very limited compared to the expressiveness of real Erlang functions.

Match specifications are given to the BIF `erlang:trace_pattern/2` to execute matching of function arguments as well as to define some actions to be taken when the match succeeds (the *MatchBody* part). Match specifications can also be used in ETS, to specify objects to be returned from an `ets:select/2` call (or other select calls). The semantics and restrictions differ slightly when using match specifications for tracing and in ETS, the differences are defined in a separate paragraph below.

The most notable difference between a *match_spec* and an Erlang fun is of course the syntax. Match specifications are Erlang terms, not Erlang code. A *match_spec* also has a somewhat strange concept of exceptions. An exception (e.g., *badarg*) in the *MatchCondition* part, which resembles an Erlang guard, will generate immediate failure, while an exception in the *MatchBody* part, which resembles the body of an Erlang function, is implicitly caught and results in the single atom `'EXIT'`.

1.1.1 Grammar

A *match_spec* can be described in this *informal* grammar:

- `MatchExpression ::= [MatchFunction, ...]`
- `MatchFunction ::= { MatchHead, MatchConditions, MatchBody }`
- `MatchHead ::= MatchVariable | '__' | [MatchHeadPart, ...]`
- `MatchHeadPart ::= term() | MatchVariable | '__'`
- `MatchVariable ::= '$<number>'`
- `MatchConditions ::= [MatchCondition, ...] | []`
- `MatchCondition ::= { GuardFunction } | { GuardFunction, ConditionExpression, ... }`

- `BoolFunction ::= is_atom | is_constant | is_float | is_integer | is_list | is_number | is_pid | is_port | is_reference | is_tuple | is_binary | is_function | is_record | is_seq_trace | 'and' | 'or' | 'not' | 'xor' | andalso | orelse`
- `ConditionExpression ::= ExprMatchVariable | { GuardFunction } | { GuardFunction, ConditionExpression, ... } | TermConstruct`
- `ExprMatchVariable ::= MatchVariable (bound in the MatchHead) | '$_' | '$$'`
- `TermConstruct = {{{}} | {{ ConditionExpression, ... }} | [] | [ConditionExpression, ...] | NonCompositeTerm | Constant`
- `NonCompositeTerm ::= term()` (not list or tuple)
- `Constant ::= {const, term() }`
- `GuardFunction ::= BoolFunction | abs | element | hd | length | node | round | size | tl | trunc | '+' | '-' | '*' | 'div' | 'rem' | 'band' | 'bor' | 'bxor' | 'bnot' | 'bsl' | 'bsr' | '>' | '>=' | '<' | '<=' | '=:=' | '===' | '=/=' | '/=' | self | get_tcw`
- `MatchBody ::= [ActionTerm]`
- `ActionTerm ::= ConditionExpression | ActionCall`
- `ActionCall ::= {ActionFunction} | {ActionFunction, ActionTerm, ...}`
- `ActionFunction ::= set_seq_token | get_seq_token | message | return_trace | process_dump | enable_trace | disable_trace | trace | display | caller | set_tcw | silent`

1.1.2 Function descriptions

Functions allowed in all types of match specifications

The different functions allowed in `match_spec` work like this:

is_atom, *is_constant*, *is_float*, *is_integer*, *is_list*, *is_number*, *is_pid*, *is_port*, *is_reference*, *is_tuple*, *is_binary*, *is_function*: Like the corresponding guard tests in Erlang, return true or false.

is_record: Takes an additional parameter, which SHALL be the result of `record_info(size, <record.type>)`, like in `{is_record, '$1', rectype, record_info(size, rectype)}`.

'not': Negates its single argument (anything other than false gives false).

'and': Returns true if all its arguments (variable length argument list) evaluate to true, else false. Evaluation order is undefined.

'or': Returns true if any of its arguments evaluates to true. Variable length argument list. Evaluation order is undefined.

andalso: Like *'and'*, but quits evaluating its arguments as soon as one argument evaluates to something else than true. Arguments are evaluated left to right.

orelse: Like *'or'*, but quits evaluating as soon as one of its arguments evaluates to true. Arguments are evaluated left to right.

'xor': Only two arguments, of which one has to be true and the other false to return true; otherwise *'xor'* returns false.

abs, *element*, *hd*, *length*, *node*, *round*, *size*, *tl*, *trunc*, *'+'*, *'-'*, *'*'*, *'div'*, *'rem'*, *'band'*, *'bor'*, *'bxor'*, *'bnot'*, *'bsl'*, *'bsr'*, *'>'*, *'>='*, *'<'*, *'<='*, *'=:='*, *'==='*, *'=/='*, *'/= '*, *self*: Work as the corresponding Erlang bif's (or operators). In case of bad arguments, the result depends on the context. In the `MatchConditions` part of the expression, the test fails immediately (like in an Erlang guard), but in the `MatchBody`, exceptions are implicitly caught and the call results in the atom `'EXIT'`.

 Functions allowed only for tracing

is_seq_trace: Returns `true` if a sequential trace token is set for the current process, otherwise `false`.

set_seq_token: Works like `seq_trace:set_token/2`, but returns `true` on success and `'EXIT'` on error or bad argument. Only allowed in the `MatchBody` part and only allowed when tracing.

get_seq_token: Works just like `seq_trace:get_token/0`, and is only allowed in the `MatchBody` part when tracing.

message: Sets an additional message appended to the trace message sent. One can only set one additional message in the body; subsequent calls will replace the appended message. As a special case, `{message, false}` disables sending of trace messages (`'call'` and `'return_to'`) for this function call, just like if the `match_spec` had not matched, which can be useful if only the side effects of the `MatchBody` are desired. Another special case is `{message, true}` which sets the default behavior, as if the function had no `match_spec`, trace message is sent with no extra information (if no other calls to `message` are placed before `{message, true}`, it is in fact a “noop”).

Takes one argument, the message. Returns `true` and can only be used in the `MatchBody` part and when tracing.

return_trace: Causes a `return_from` trace message to be sent upon return from the current function. Takes no arguments, returns `true` and can only be used in the `MatchBody` part when tracing. If the process trace flag `silent` is active the `return_from` trace message is inhibited.

NOTE! If the traced function is tail recursive, this match spec function destroys that property. Hence, if a match spec executing this function is used on a perpetual server process, it may only be active for a limited time, or the emulator will eventually use all memory in the host machine and crash. If this `match_spec` function is inhibited using the `silent` process trace flag tail recursiveness still remains.

exception_trace: Same as *return_trace*, plus; if the traced function exits due to an exception, an `exception_from` trace message is generated, whether the exception is caught or not.

process_dump: Returns some textual information about the current process as a binary. Takes no arguments and is only allowed in the `MatchBody` part when tracing.

enable_trace: With one parameter this function turns on tracing like the Erlang call `erlang:trace(self(), true, [P2])`, where `P2` is the parameter to `enable_trace`. With two parameters, the first parameter should be either a process identifier or the registered name of a process. In this case tracing is turned on for the designated process in the same way as in the Erlang call `erlang:trace(P1, true, [P2])`, where `P1` is the first and `P2` is the second argument. The process `P1` gets its trace messages sent to the same tracer as the process executing the statement uses. `P1` can *not* be one of the atoms `all`, `new` or `existing` (unless, of course, they are registered names). `P2` can *not* be `cpu_timestamp` nor `{tracer, _}`. Returns `true` and may only be used in the `MatchBody` part when tracing.

disable_trace: With one parameter this function disables tracing like the Erlang call `erlang:trace(self(), false, [P2])`, where `P2` is the parameter to `disable_trace`. With two parameters it works like the Erlang call `erlang:trace(P1, false, [P2])`, where `P1` can be either a process identifier or a registered name and is given as the first argument to the `match_spec` function. `P2` can *not* be `cpu_timestamp` nor `{tracer, _}`. Returns `true` and may only be used in the `MatchBody` part when tracing.

trace: With two parameters this function takes a list of trace flags to disable as first parameter and a list of trace flags to enable as second parameter. Logically, the disable list is applied first, but effectively all changes are applied atomically. The trace flags are the same as for `erlang:trace/3` not including `cpu_timestamp` but including `{tracer, _}`. If a tracer is specified in both lists, the tracer in the enable list takes precedence. If no tracer is specified the same tracer as the process executing the match spec is used. With three parameters to this function the first is either a process identifier or the registered name of a process to set trace flags on, the second is the disable list, and the third is the enable list. Returns

`true` if any trace property was changed for the trace target process or `false` if not. It may only be used in the `MatchBody` part when tracing.

caller: Returns the calling function as a tuple `{Module, Function, Arity}` or the atom `undefined` if the calling function cannot be determined. May only be used in the `MatchBody` part when tracing.

Note that if a “technically built in function” (i.e. a function not written in Erlang) is traced, the `caller` function will sometimes return the atom `undefined`. The calling Erlang function is not available during such calls.

display: For debugging purposes only; displays the single argument as an Erlang term on stdout, which is seldom what is wanted. Returns `true` and may only be used in the `MatchBody` part when tracing.

get_tcw: Takes no argument and returns the value of the node's trace control word. The same is done by `erlang:system_info(trace_control_word)`.

The trace control word is a 32-bit unsigned integer intended for generic trace control. The trace control word can be tested and set both from within trace match specifications and with BIFs. This call is only allowed when tracing.

set_tcw: Takes one unsigned integer argument, sets the value of the node's trace control word to the value of the argument and returns the previous value. The same is done by `erlang:system_flag(trace_control_word, Value)`. It is only allowed to use `set_tcw` in the `MatchBody` part when tracing.

silent: Takes one argument. If the argument is `true`, the call trace message mode for the current process is set to silent for this call and all subsequent, i.e. call trace messages are inhibited even if `{message, true}` is called in the `MatchBody` part for a traced function.

This mode can also be activated with the `silent` flag to `erlang:trace/3`.

If the argument is `false`, the call trace message mode for the current process is set to normal (non-silent) for this call and all subsequent.

If the argument is neither `true` nor `false`, the call trace message mode is unaffected.

Note that all “function calls” have to be tuples, even if they take no arguments. The value of `self` is the atom `self`, but the value of `{self}` is the `pid()` of the current process.

1.1.3 Variables and literals

Variables take the form `'$<number>'` where `<number>` is an integer between 0 (zero) and 100000000 (`1e+8`), the behavior if the number is outside these limits is *undefined*. In the `MatchHead` part, the special variable `'_'` matches anything, and never gets bound (like `_` in Erlang). In the `MatchCondition/MatchBody` parts, no unbound variables are allowed, why `'_'` is interpreted as itself (an atom). Variables can only be bound in the `MatchHead` part. In the `MatchBody` and `MatchCondition` parts, only variables bound previously may be used. As a special case, in the `MatchCondition/MatchBody` parts, the variable `'$_'` expands to the whole expression which matched the `MatchHead` (i.e., the whole parameter list to the possibly traced function or the whole matching object in the ets table) and the variable `'$$'` expands to a list of the values of all bound variables in order (i.e. `['$1', '$2', ...]`).

In the `MatchHead` part, all literals (except the variables noted above) are interpreted as is. In the `MatchCondition/MatchBody` parts, however, the interpretation is in some ways different. Literals in the `MatchCondition/MatchBody` can either be written as is, which works for all literals except tuples, or by using the special form `{const, T}`, where `T` is any Erlang term. For tuple literals in the `match_spec`, one can also use double tuple parentheses, i.e., construct them as a tuple of arity one containing a single tuple, which is the one to be constructed. The “double tuple parenthesis” syntax is useful to construct tuples from already bound variables, like in `{{'$1', [a, b, '$2']}`. Some examples may be needed:

Expression	Variable bindings	Result
<code>{{'S1','S2'}}</code>	<code>'S1' = a, 'S2' = b</code>	<code>{a,b}</code>
<code>{const, {'S1','S2'}}</code>	doesn't matter	<code>{'S1','S2'}</code>
<code>a</code>	doesn't matter	<code>a</code>
<code>'S1'</code>	<code>'S1' = []</code>	<code>[]</code>
<code>['S1']</code>	<code>'S1' = []</code>	<code>[[]]</code>
<code>{{{a}}}</code>	doesn't matter	<code>{{a}}</code>
<code>42</code>	doesn't matter	<code>42</code>
<code>"hello"</code>	doesn't matter	<code>"hello"</code>
<code>\$1</code>	doesn't matter	49 (the ASCII value for the character '1')

Table 1.1: Literals in the MatchCondition/MatchBody parts of a match_spec

1.1.4 Execution of the match

The execution of the match expression, when the runtime system decides whether a trace message should be sent, goes as follows:

For each tuple in the MatchExpression list and while no match has succeeded:

- Match the MatchHead part against the arguments to the function, binding the '`$<number>`' variables (much like in `ets:match/2`). If the MatchHead cannot match the arguments, the match fails.
- Evaluate each MatchCondition (where only '`$<number>`' variables previously bound in the MatchHead can occur) and expect it to return the atom `true`. As soon as a condition does not evaluate to `true`, the match fails. If any BIF call generates an exception, also fail.
- - *If the match_spec is executing when tracing:*
Evaluate each ActionTerm in the same way as the MatchConditions, but completely ignore the return values. Regardless of what happens in this part, the match has succeeded.
 - *If the match_spec is executed when selecting objects from an ETS table:*
Evaluate the expressions in order and return the value of the last expression (typically there is only one expression in this context)

1.1.5 Differences between match specifications in ETS and tracing

ETS match specifications are there to produce a return value. Usually the expression contains one single ActionTerm which defines the return value without having any side effects. Calls with side effects are not allowed in the ETS context.

When tracing there is no return value to produce, the match specification either matches or doesn't. The effect when the expression matches is a trace message rather than a returned term. The ActionTerm's are executed as in an imperative language, i.e. for their side effects. Functions with side effects are also allowed when tracing.

In ETS the match head is a `tuple()` (or a single match variable) while it is a list (or a single match variable) when tracing.

1.1.6 Examples

Match an argument list of three where the first and third arguments are equal:

```
[{'$1', '_ ', '$1'},
 [],
 []]
```

Match an argument list of three where the second argument is a number greater than three:

```
[['_ ', '$1', '_ '],
 [{'>', '$1', 3}],
 []]
```

Match an argument list of three, where the third argument is a tuple containing argument one and two or a list beginning with argument one and two (i. e. [a,b,[a,b,c]] or [a,b,{a,b}]):

```
[{'$1', '$2', '$3'},
 [{orelse,
   {':=', '$3', {'$1','$2'}},
   {'and',
    {':=', '$1', {hd, '$3'}},
    {':=', '$2', {hd, {tl, '$3'}}}}}],
 []]
```

The above problem may also be solved like this:

```
[{'$1', '$2', {'$1', '$2'}], [], [],
 [{'$1', '$2', ['$1', '$2' | '_ ']}, [], []]
```

Match two arguments where the first is a tuple beginning with a list which in turn begins with the second argument times two (i. e. [{[4,x],y},2] or [{[8], y, z},4])

```
[{'$1', '$2'}, [{':=', {'*', 2, '$2'}, {hd, {element, 1, '$1'}}}],
 []]
```

Match three arguments. When all three are equal and are numbers, append the process dump to the trace message, else let the trace message be as is, but set the sequential trace token label to 4711.

```
[{'$1', '$1', '$1'},
 [{is_number, '$1'}],
 [{message, {process_dump}}}],
 {'_', [], [{set_seq_token, label, 4711}]}
```

As can be noted above, the parameter list can be matched against a single MatchVariable or an '_ '. To replace the whole parameter list with a single variable is a special case. In all other cases the MatchHead has to be a *proper* list.

Match all objects in an ets table where the first element is the atom 'strider' and the tuple arity is 3 and return the whole object.

```
[{{strider, '_'. '_'},
 [],
 ['$_']}]
```

Match all objects in an ets table with arity > 1 and the first element is 'gandalf', return element 2.

```
[{'$1',
  [{'==' , gandalf, {element, 1, '$1'}}, {'>=' , {size, '$1'}, 2}],
  [{element, 2, '$1'}]]
```

In the above example, if the first element had been the key, it's much more efficient to match that key in the `MatchHead` part than in the `MatchConditions` part. The search space of the tables is restricted with regards to the `MatchHead` so that only objects with the matching key are searched.

Match tuples of 3 elements where the second element is either 'merry' or 'pippin', return the whole objects.

```
[{'_ ', merry, '_ '},
 [],
 ['$_ '}],
 {'_ ', pippin, '_ '},
 [],
 ['$_ ']]
```

The function `ets:test_ms/2` can be useful for testing complicated ets matches.

1.2 How to interpret the Erlang crash dumps

This document describes the `erl_crash.dump` file generated upon abnormal exit of the Erlang runtime system.

Important: For OTP release R9C the Erlang crash dump has had a major facelift. This means that the information in this document will not be directly applicable for older dumps. However, if you use the `Crashdump Viewer` tool on older dumps, the crash dumps are translated into a format similar to this.

The system will write the crash dump in the current directory of the emulator or in the file pointed out by the environment variable (whatever that means on the current operating system)

`ERL_CRASH_DUMP`. For a crash dump to be written, there has to be a writable file system mounted.

Crash dumps are written mainly for one of two reasons: either the builtin function `erlang:halt/1` is called explicitly with a string argument from running Erlang code, or else the runtime system has detected an error that cannot be handled. The most usual reason that the system can't handle the error is that the cause is external limitations, such as running out of memory. A crash dump due to an internal error may be caused by the system reaching limits in the emulator itself (like the number of atoms in the system, or too many simultaneous ets tables). Usually the emulator or the operating system can be reconfigured to avoid the crash, which is why interpreting the crash dump correctly is important.

The erlang crash dump is a readable text file, but it might not be very easy to read. Using the `Crashdump Viewer` tool in the `observer` application will simplify the task. This is an HTML based tool for browsing Erlang crash dumps.

1.2.1 General information

The first part of the dump shows the creation time for the dump, a slogan indicating the reason for the dump, the system version, of the node from which the dump originates, the compile time of the emulator running the originating node and the number of atoms in the atom table.

Reasons for crash dumps (slogan)

The reason for the dump is noted in the beginning of the file as *Slogan: <reason>* (the word “slogan” has historical roots). If the system is halted by the BIF `erlang:halt/1`, the slogan is the string parameter passed to the BIF, otherwise it is a description generated by the emulator or the (Erlang) kernel.

Normally the message should be enough to understand the problem, but nevertheless some messages are described here. Note however that the suggested reasons for the crash are *only suggestions*. The exact reasons for the errors may vary depending on the local applications and the underlying operating system.

- “<A>: Cannot allocate <N> bytes of memory (of type “<T>”).” - The system has run out of memory. <A> is the allocator that failed to allocate memory, <N> is the number of bytes that <A> tried to allocate, and <T> is the memory block type that the memory was needed for. The most common case is that a process stores huge amounts of data. In this case <T> is most often `heap`, `old_heap`, `heap_frag`, or `binary`. For more information on allocators see `erts_alloc(3)` [page 104].
- “<A>: Cannot reallocate <N> bytes of memory (of type “<T>”).” - Same as above with the exception that memory was being reallocated instead of being allocated when the system ran out of memory.
- “Unexpected op code N” - Error in compiled code, beam file damaged or error in the compiler.
- “Module *Name* undefined” | “Function *Name* undefined” | “No function *Name:Name/1*” | “No function *Name:start/2*” - The kernel/stdlib applications are damaged or the start script is damaged.
- “Driver_select called with too large file descriptor N” - The number of file descriptors for sockets exceed 1024 (Unix only). The limit on file-descriptors in some Unix flavors can be set to over 1024, but only 1024 sockets/pipes can be used simultaneously by Erlang (due to limitations in the Unix `select` call). The number of open regular files is not affected by this.
- “Received SIGUSR1” - The SIGUSR1 signal was sent to the Erlang machine (Unix only).
- “Kernel pid terminated (*Who*) (*Exit-reason*)” - The kernel supervisor has detected a failure, usually that the `application_controller` has shut down (`Who = application_controller`, `Why = shutdown`). The application controller may have shut down for a number of reasons, the most usual being that the node name of the distributed Erlang node is already in use. A complete supervisor tree “crash” (i.e., the top supervisors have exited) will give about the same result. This message comes from the Erlang code and not from the virtual machine itself. It is always due to some kind of failure in an application, either within OTP or a “user-written” one. Looking at the error log for your application is probably the first step to take.
- “Init terminating in `do_boot ()`” - The primitive Erlang boot sequence was terminated, most probably because the boot script has errors or cannot be read. This is usually a configuration error - the system may have been started with a faulty `-boot` parameter or with a boot script from the wrong version of OTP.
- “Could not start kernel pid (*Who*) ()” - One of the kernel processes could not start. This is probably due to faulty arguments (like errors in a `-config` argument) or faulty configuration files. Check that all files are in their correct location and that the configuration files (if any) are not damaged. Usually there are also messages written to the controlling terminal and/or the error log explaining what’s wrong.

Other errors than the ones mentioned above may occur, as the `erlang:halt/1` BIF may generate any message. If the message is not generated by the BIF and does not occur in the list above, it may be due to an error in the emulator. There may however be unusual messages that I haven't mentioned, that still are connected to an application failure. There is a lot more information available, so more thorough reading of the crash dump may reveal the crash reason. The size of processes, the number of ets tables and the Erlang data on each process stack can be useful for tracking down the problem.

Number of atoms

The number of atoms in the system at the time of the crash is shown as *Atoms: <number>*. Some ten thousands atoms is perfectly normal, but more could indicate that the BIF `erlang:list_to_atom/1` is used to dynamically generate a lot of *different* atoms, which is never a good idea.

1.2.2 Memory information

Under the tag `=memory` you will find information similar to what you can obtain on a living node with `[erlang:memory0]`.

1.2.3 Internal table information

The tags `=hash_table:<table_name>` and `=index_table:<table_name>` presents internal tables. These are mostly of interest for runtime system developers.

1.2.4 Allocated areas

Under the tag `=allocated_areas` you will find information similar to what you can obtain on a living node with `[erlang:system_info(allocated_areas)]`.

1.2.5 Allocator

Under the tag `=allocator:<A>` you will find various information about allocator `<A>`. The information is similar to what you can obtain on a living node with `[erlang:system_info({allocator, <A>})]`. For more information see the documentation of `[erlang:system_info({allocator, <A>})]`, and the `erts_alloc(3)` [page 104] documentation.

1.2.6 Process information

The Erlang crashdump contains a listing of each living Erlang process in the system. The process information for one process may look like this (line numbers have been added):

The following fields can exist for a process:

`=proc:<pid>` Heading, states the process identifier

State The state of the process. This can be one of the following:

- *Scheduled* - The process was scheduled to run but not currently running (“in the run queue”).
- *Waiting* - The process was waiting for something (in receive).
- *Running* - The process was currently running. If the BIF `erlang:halt/1` was called, this was the process calling it.
- *Exiting* - The process was on its way to exit.

- *Garbing* - This is bad luck, the process was garbage collecting when the crash dump was written, the rest of the information for this process is limited.
- *Suspended* - The process is suspended, either by the BIF `erlang:suspend_process/1` or because it is trying to write to a busy port.

Registered name The registered name of the process, if any.

Spawned as The entry point of the process, i.e., what function was referenced in the `spawn` or `spawn_link` call that started the process.

Last scheduled in for | **Current call** The current function of the process. These fields will not always exist.

Spawned by The parent of the process, i.e. the process which executed `spawn` or `spawn_link`.

Started The date and time when the process was started.

Message queue length The number of messages in the process' message queue.

Number of heap fragments The number of allocated heap fragments.

Heap fragment data Size of fragmented heap data. This is data either created by messages being sent to the process or by the Erlang BIFs. This amount depends on so many things that this field is utterly uninteresting.

Link list Process id's of processes linked to this one. May also contain ports. If process monitoring is used, this field also tells in which direction the monitoring is in effect, i.e., a link being "to" a process tells you that the "current" process was monitoring the other and a link "from" a process tells you that the other process was monitoring the current one.

Reductions The number of reductions consumed by the process.

Stack+heap The size of the stack and heap (they share memory segment)

OldHeap The size of the "old heap". The Erlang virtual machine uses generational garbage collection with two generations. There is one heap for new data items and one for the data that have survived two garbage collections. The assumption (which is almost always correct) is that data that survive two garbage collections can be "tenured" to a heap more seldom garbage collected, as they will live for a long period. This is a quite usual technique in virtual machines. The sum of the heaps and stack together constitute most of the process's allocated memory.

Heap unused, OldHeap unused The amount of unused memory on each heap. This information is usually useless.

Stack If the system uses shared heap, the fields *Stack+heap*, *OldHeap*, *Heap unused* and *OldHeap unused* do not exist. Instead this field presents the size of the process' stack.

Program counter The current instruction pointer. This is only interesting for runtime system developers. The function into which the program counter points is the current function of the process.

CP The continuation pointer, i.e. the return address for the current call. Usually useless for other than runtime system developers. This may be followed by the function into which the CP points, which is the function calling the current function.

Arity The number of live argument registers. The argument registers, if any are live, will follow. These may contain the arguments of the function if they are not yet moved to the stack.

See also the section about process data [page 12].

1.2.7 Port information

This section lists the open ports, their owners, any linked processes, and the name of their driver or external process.

1.2.8 ETS tables

This section contains information about all the ETS tables in the system. The following fields are interesting for each table:

=ets:<owner> Heading, states the owner of the table (a process identifier)

Table The identifier for the table. If the table is a `named_table`, this is the name.

Name The name of the table, regardless of whether it is a `named_table` or not.

Buckets This occurs if the table is a hash table, i.e. if it is not an `ordered_set`.

Ordered set (AVL tree), Elements This occurs only if the table is an `ordered_set`. (The number of elements is the same as the number of objects in the table.)

Objects The number of objects in the table

Words The number of words (usually 4 bytes/word) allocated to data in the table.

1.2.9 Timers

This section contains information about all the timers started with the BIFs `erlang:start_timer/3` and `erlang:send_after/3`. The following fields exist for each timer:

=timer:<owner> Heading, states the owner of the timer (a process identifier) i.e. the process to receive the message when the timer expires.

Message The message to be sent.

Time left Number of milliseconds left until the message would have been sent.

1.2.10 Distribution information

If the Erlang node was alive, i.e., set up for communicating with other nodes, this section lists the connections that were active. The following fields can exist:

=node:<node_name> The name of the node

no-distribution This will only occur if the node was not distributed.

=visible_node:<channel> Heading for a visible node, i.e. an alive node with a connection to the node that crashed. States the channel number for the node.

=hidden_node:<channel> Heading for a hidden node. A hidden node is the same as a visible node, except that it is started with the “-hidden” flag. States the channel number for the node.

=not_connected:<channel> Heading for a node which has been connected to the crashed node earlier. References (i.e. process or port identifiers) to the not connected node existed at the time of the crash. exist. States the channel number for the node.

Name The name of the remote node.

Controller The port which controls the communication with the remote node.

Creation An integer (1-3) which together with the node name identifies a specific instance of the node.

Remote monitoring: <local_proc> <remote_proc> The local process was monitoring the remote process at the time of the crash.

Remotely monitored by: <local_proc> <remote_proc> The remote process was monitoring the local process at the time of the crash.

Remote link: <local_proc> <remote_proc> A link existed between the local process and the remote process at the time of the crash.

1.2.11 Loaded module information

This section contains information about all loaded modules. First, the memory usage by loaded code is summarized. There is one field for “Current code” which is code that is the current latest version of the modules. There is also a field for “Old code” which is code where there exists a newer version in the system, but the old version is not yet purged. The memory usage is in bytes.

All loaded modules are then listed. The following fields exist:

=mod:*<module_name>* Heading, and the name of the module.

Current size Memory usage for the loaded code in bytes

Old size Memory usage for the old code, if any.

Current attributes Module attributes for the current code. This field is decoded when looked at by the Crashdump Viewer tool.

Old attributes Module attributes for the old code, if any. This field is decoded when looked at by the Crashdump Viewer tool.

Current compilation info Compilation information (options) for the current code. This field is decoded when looked at by the Crashdump Viewer tool.

Old compilation info Compilation information (options) for the old code, if any. This field is decoded when looked at by the Crashdump Viewer tool.

1.2.12 Fun information

In this section, all funs are listed. The following fields exist for each fun:

=fun Heading

Module The name of the module where the fun was defined.

Uniq, Index Identifiers

Address The address of the fun's code.

Native_address The address of the fun's code when HiPE is enabled.

Refc The number of references to the fun.

1.2.13 Process Data

For each process there will be at least one *=proc_stack* and one *=proc_heap* tag followed by the raw memory information for the stack and heap of the process.

For each process there will also be a *=proc_messages* tag if the process' message queue is non-empty and a *=proc_dictionary* tag if the process' dictionary (the put/2 and get/1 thing) is non-empty.

The raw memory information can be decoded by the Crashdump Viewer tool. You will then be able to see the stack dump, the message queue (if any) and the dictionary (if any).

The stack dump is a dump of the Erlang process stack. Most of the live data (i.e., variables currently in use) are placed on the stack; thus this can be quite interesting. One has to “guess” what's what, but as the information is symbolic, thorough reading of this information can be very useful. As an example we can find the state variable of the Erlang primitive loader on line (5) in the example below:

```
(1) 3cac44   Return addr 0x13BF58 (<terminate process normally>)
(2) y(0)     ["/view/siri_r10_dev/clearcase/otp/erts/lib/kernel/ebin", "/view/siri_r10_dev/
(3) clearcase/otp/erts/lib/stdlib/ebin"]
(4) y(1)     <0.1.0>
(5) y(2)     {state, [], none, #Fun<erl_prim_loader.6.7085890>, undefined, #Fun<erl_prim_loader.7.9000
(6) y(3)     infinity
```

When interpreting the data for a process, it is helpful to know that anonymous function objects (funs) are given a name constructed from the name of the function in which they are created, and a number (starting with 0) indicating the number of that fun within that function.

1.2.14 Atoms

Now all the atoms in the system are written. This is only interesting if one suspects that dynamic generation of atoms could be a problem, otherwise this section can be ignored.

Note that the last created atom is printed first.

1.2.15 Disclaimer

The format of the crash dump evolves between releases of OTP. Some information here may not apply to your version. A description as this will never be complete; it is meant as an explanation of the crash dump in general and as a help when trying to find application errors, not as a complete specification.

1.3 How to implement an alternative carrier for the Erlang distribution

This document describes how one can implement ones own carrier protocol for the Erlang distribution. The distribution is normally carried by the TCP/IP protocol. What's explained here is the method for replacing TCP/IP with another protocol.

The document is a step by step explanation of the `uds_dist` example application (seated in the kernel applications `examples` directory). The `uds_dist` application implements distribution over Unix domain sockets and is written for the Sun Solaris 2 operating environment. The mechanisms are however general and applies to any operating system Erlang runs on. The reason the C code is not made portable, is simply readability.

1.3.1 Introduction

To implement a new carrier for the Erlang distribution, one must first make the protocol available to the Erlang machine, which involves writing an Erlang driver. There is no way one can use a port program, there *has* to be an Erlang driver. Erlang drivers can either be statically linked to the emulator, which can be an alternative when using the open source distribution of Erlang, or dynamically loaded into the Erlang machines address space, which is the only alternative if a precompiled version of Erlang is to be used.

Writing an Erlang driver is by no means easy. The driver is written as a couple of call-back functions called by the Erlang emulator when data is sent to the driver or the driver has any data available on a file descriptor. As the driver call-back routines execute in the main thread of the Erlang machine, the call-back functions can perform no blocking activity whatsoever. The call-backs should only set up file descriptors for waiting and/or read/write available data. All I/O has to be non blocking. Driver

call-backs are however executed in sequence, why a global state can safely be updated within the routines.

When the driver is implemented, one would preferably write an Erlang interface for the driver to be able to test the functionality of the driver separately. This interface can then be used by the distribution module which will cover the details of the protocol from the `net_kernel`. The easiest path is to mimic the `inet` and `gen_tcp` interfaces, but a lot of functionality in those modules need not be implemented. In the example application, only a few of the usual interfaces are implemented, and they are much simplified.

When the protocol is available to Erlang through a driver and an Erlang interface module, a distribution module can be written. The distribution module is a module with well defined call-backs, much like a `gen_server` (there is no compiler support for checking the call-backs though). The details of finding other nodes (i.e. talking to `epmd` or something similar), creating a listen port (or similar), connecting to other nodes and performing the handshakes/cookie verification are all implemented by this module. There is however a utility module, `dist_util`, that will do most of the hard work of handling handshakes, cookies, timers and ticking. Using `dist_util` makes implementing a distribution module much easier and that's what we are doing in the example application.

The last step is to create boot scripts to make the protocol implementation available at boot time. The implementation can be debugged by starting the distribution when all of the system is running, but in a real system the distribution should start very early, why a boot-script and some command line parameters are necessary. This last step also implies that the Erlang code in the interface and distribution modules is written in such a way that it can be run in the startup phase. Most notably there can be no calls to the application module or to any modules not loaded at boot-time (i.e. only `kernel`, `stdlib` and the application itself can be used).

1.3.2 The driver

Although Erlang drivers in general may be beyond the scope of this document, a brief introduction seems to be in place.

Drivers in general

An Erlang driver is a native code module written in C (or assembler) which serves as an interface for some special operating system service. This is a general mechanism that is used throughout the Erlang emulator for all kinds of I/O. An Erlang driver can be dynamically linked (or loaded) to the Erlang emulator at runtime by using the `erl_ddll` Erlang module. Some of the drivers in OTP are however statically linked to the runtime system, but that's more an optimization than a necessity.

The driver data-types and the functions available to the driver writer are defined in the header file `erl_driver.h` (there is also an deprecated version called `driver.h`, don't use that one.) seated in Erlang's include directory (and in `$ERL_TOP/erts/emulator/beam` in the source code distribution). Refer to that file for function prototypes etc.

When writing a driver to make a communications protocol available to Erlang, one should know just about everything worth knowing about that particular protocol. All operation has to be non blocking and all possible situations should be accounted for in the driver. A non stable driver will affect and/or crash the whole Erlang runtime system, which is seldom what's wanted.

The emulator calls the driver in the following situations:

- When the driver is loaded. This call-back has to have a special name and will inform the emulator of what call-backs should be used by returning a pointer to a `Er1DrvEntry` struct, which should be properly filled in (see below).

- When a port to the driver is opened (by a `open_port` call from Erlang). This routine should set up internal data structures and return an opaque data entity of the type `Er1DrvData`, which is a data-type large enough to hold a pointer. The pointer returned by this function will be the first argument to all other call-backs concerning this particular port. It is usually called the port handle. The emulator only stores the handle and does never try to interpret it, why it can be virtually anything (well anything not larger than a pointer that is) and can point to anything if it is a pointer. Usually this pointer will refer to a structure holding information about the particular port, as it does in our example.
- When an Erlang process sends data to the port. The data will arrive as a buffer of bytes, the interpretation is not defined, but is up to the implementor. This call-back returns nothing to the caller, answers are sent to the caller as messages (using a routine called `driver_output` available to all drivers). There is also a way to talk in a synchronous way to drivers, described below. There can be an additional call-back function for handling data that is fragmented (sent in a deep io-list). That interface will get the data in a form suitable for Unix `writew` rather than in a single buffer. There is no need for a distribution driver to implement such a call-back, so we wont.
- When a file descriptor is signaled for input. This call-back is called when the emulator detects input on a file descriptor which the driver has marked for monitoring by using the interface `driver_select`. The mechanism of driver select makes it possible to read non blocking from file descriptors by calling `driver_select` when reading is needed and then do the actual reading in this call-back (when reading is actually possible). The typical scenario is that `driver_select` is called when an Erlang process orders a read operation, and that this routine sends the answer when data is available on the file descriptor.
- When a file descriptor is signaled for output. This call-back is called in a similar way as the previous, but when writing to a file descriptor is possible. The usual scenario is that Erlang orders writing on a file descriptor and that the driver calls `driver_select`. When the descriptor is ready for output, this call-back is called and the driver can try to send the output. There may of course be queuing involved in such operations, and there are some convenient queue routines available to the driver writer to use in such situations.
- When a port is closed, either by an Erlang process or by the driver calling one of the `driver_failure_XXX` routines. This routine should clean up everything connected to one particular port. Note that when other call-backs call a `driver_failure_XXX` routine, this routine will be immediately called and the call-back routine issuing the error can make no more use of the data structures for the port, as this routine surely has freed all associated data and closed all file descriptors. If the queue utility available to driver writes is used, this routine will however *not* be called until the queue is empty.
- When an Erlang process calls `erlang:driver_control/2`, which is a synchronous interface to drivers. The control interface is used to set driver options, change states of ports etc. We'll use this interface quite a lot in our example.
- When a timer expires. The driver can set timers with the function `driver_set_timer`. When such timers expire, a specific call-back function is called. We will not use timers in our example.
- When the whole driver is unloaded. Every resource allocated by the driver should be freed.

The distribution driver's data structures

The driver used for Erlang distribution should implement a reliable, order maintaining, variable length packet oriented protocol. All error correction, re-sending and such need to be implemented in the driver or by the underlying communications protocol. If the protocol is stream oriented (as is the case with both TCP/IP and our streamed Unix domain sockets), some mechanism for packaging is needed. We will use the simple method of having a header of four bytes containing the length of the package in a big endian 32 bit integer (as Unix domain sockets only can be used between processes on the same

machine, we actually don't need to code the integer in some special endianness, but I'll do it anyway because in most situation you do need to do it. Unix domain sockets are reliable and order maintaining, so we don't need to implement resends and such in our driver.

Lets start writing our example Unix domain sockets driver by declaring prototypes and filling in a static ErlDrvEntry structure.

```
( 1) #include <stdio.h>
( 2) #include <stdlib.h>
( 3) #include <string.h>
( 4) #include <unistd.h>
( 5) #include <errno.h>
( 6) #include <sys/types.h>
( 7) #include <sys/stat.h>
( 8) #include <sys/socket.h>
( 9) #include <sys/un.h>
(10) #include <fcntl.h>

(11) #define HAVE_UIO_H
(12) #include "erl_driver.h"

(13) /*
(14) ** Interface routines
(15) */
(16) static ErlDrvData uds_start(ErlDrvPort port, char *buff);
(17) static void uds_stop(ErlDrvData handle);
(18) static void uds_command(ErlDrvData handle, char *buff, int buflen);
(19) static void uds_input(ErlDrvData handle, ErlDrvEvent event);
(20) static void uds_output(ErlDrvData handle, ErlDrvEvent event);
(21) static void uds_finish(void);
(22) static int uds_control(ErlDrvData handle, unsigned int command,
(23)                       char* buf, int count, char** res, int res_size);

(24) /* The driver entry */
(25) static ErlDrvEntry uds_driver_entry = {
(26)     NULL, /* init, N/A */
(27)     uds_start, /* start, called when port is opened */
(28)     uds_stop, /* stop, called when port is closed */
(29)     uds_command, /* output, called when erlang has sent */
(30)     uds_input, /* ready_input, called when input descriptor
(31)               ready */
(32)     uds_output, /* ready_output, called when output
(33)               descriptor ready */
(34)     "uds_drv", /* char *driver_name, the argument
(35)               to open_port */
(36)     uds_finish, /* finish, called when unloaded */
(37)     NULL, /* void * that is not used (BC) */
(38)     uds_control, /* control, port_control callback */
(39)     NULL, /* timeout, called on timeouts */
(40)     NULL /* outputv, vector output interface */
(41) };
```

On line 1 to 10 we have included the OS headers needed for our driver. As this driver is written for Solaris, we know that the header `uio.h` exists, why we can define the preprocessor variable `HAVE_UIO_H`

before we include `erl_driver.h` at line 12. The definition of `HAVE_UIO_H` will make the I/O vectors used in Erlang's driver queues to correspond to the operating systems ditto, which is very convenient.

The different call-back functions are declared ("forward declarations") on line 16 to 23.

The driver structure is similar for statically linked in drivers and dynamically loaded. However some of the fields should be left empty (i.e. initialized to `NULL`) in the different types of drivers. The first field (the `init` function pointer) is always left blank in a dynamically loaded driver, which can be seen on line 26. The `NULL` on line 37 should always be there, the field is no longer used and is retained for backward compatibility. We use no timers in this driver, why no call-back for timers is needed. The last field (line 40) can be used to implement an interface similar to Unix `writew` for output. There is no need for such interface in a distribution driver, so we leave it with a `NULL` value (We will however use scatter/gather I/O internally in the driver).

Our defined call-backs thus are:

- `uds_start`, which shall initiate data for a port. We won't create any actual sockets here, just initialize data structures.
- `uds_stop`, the function called when a port is closed.
- `uds_command`, which will handle messages from Erlang. The messages can either be plain data to be sent or more subtle instructions to the driver. We will use this function mostly for data pumping.
- `uds_input`, this is the call-back which is called when we have something to read from a socket.
- `uds_output`, this is the function called when we can write to a socket.
- `uds_finish`, which is called when the driver is unloaded. A distribution driver will actually (or hopefully) never be unloaded, but we include this for completeness. Being able to clean up after oneself is always a good thing.
- `uds_control`, the `erlang:port_control/2` call-back, which will be used a lot in this implementation.

The ports implemented by this driver will operate in two major modes, which I will call the *command* and *data* modes. In command mode, only passive reading and writing (like `gen_tcp:recv/gen_tcp:send`) can be done, and this is the mode the port will be in during the distribution handshake. When the connection is up, the port will be switched to data mode and all data will be immediately read and passed further to the Erlang emulator. In data mode, no data arriving to the `uds_command` will be interpreted, but just packaged and sent out on the socket. The `uds_control` call-back will do the switching between those two modes.

While the `net_kernel` informs different subsystems that the connection is coming up, the port should accept data to send, but not receive any data, to avoid that data arrives from another node before every kernel subsystem is prepared to handle it. We have a third mode for this intermediate stage, let's call it the *intermediate* mode.

Let's define an enum for the different types of ports we have:

```
( 1) typedef enum {
( 2)     portTypeUnknown,      /* An uninitialized port */
( 3)     portTypeListener,    /* A listening port/socket */
( 4)     portTypeAcceptor,    /* An intermediate stage when accepting
( 5)                               on a listen port */
( 6)     portTypeConnector,   /* An intermediate stage when connecting */
( 7)     portTypeCommand,     /* A connected open port in command mode */
( 8)     portTypeIntermediate, /* A connected open port in special
( 9)                               half active mode */
(10)    portTypeData         /* A connected open port in data mode */
```

```
(11) } PortType;
```

Lets look at the different types:

- portTypeUnknown - The type a port has when it's opened, but not actually bound to any file descriptor.
- portTypeListener - A port that is connected to a listen socket. This port will not do especially much, there will be no data pumping done on this socket, but there will be read data available when one is trying to do an accept on the port.
- portTypeAcceptor - This is a port that is to represent the result of an accept operation. It is created when one wants to accept from a listen socket, and it will be converted to a portTypeCommand when the accept succeeds.
- portTypeConnector - Very similar to portTypeAcceptor, an intermediate stage between the request for a connect operation and that the socket is really connected to an accepting ditto in the other end. As soon as the sockets are connected, the port will switch type to portTypeCommand.
- portTypeCommand - A connected socket (or accepted socket if you want) that is in the command mode mentioned earlier.
- portTypeIntermediate - The intermediate stage for a connected socket. There should be no processing of input for this socket.
- portTypeData - The mode where data is pumped through the port and the uds_command routine will regard every call as a call where sending is wanted. In this mode all input available will be read and sent to Erlang as soon as it arrives on the socket, much like in the active mode of a gen_tcp socket.

Now lets look at the state we'll need for our ports. One can note that not all fields are used for all types of ports and that one could save some space by using unions, but that would clutter the code with multiple indirections, so i simply use one struct for all types of ports, for readability.

```
( 1) typedef unsigned char Byte;
( 2) typedef unsigned int Word;

( 3) typedef struct uds_data {
( 4)     int fd;                /* File descriptor */
( 5)     ErlDrvPort port;      /* The port identifier */
( 6)     int lockfd;          /* The file descriptor for a lock file in
( 7)                          case of listen sockets */
( 8)     Byte creation;        /* The creation serial derived from the
( 9)                          lockfile */
(10)     PortType type;        /* Type of port */
(11)     char *name;          /* Short name of socket for unlink */
(12)     Word sent;           /* Bytes sent */
(13)     Word received;       /* Bytes received */
(14)     struct uds_data *partner; /* The partner in an accept/listen pair */
(15)     struct uds_data *next; /* Next structure in list */
(16)     /* The input buffer and it's data */
(17)     int buffer_size;      /* The allocated size of the input buffer */
(18)     int buffer_pos;      /* Current position in input buffer */
(19)     int header_pos;      /* Where the current header is in the
(20)                          input buffer */
(21)     Byte *buffer;        /* The actual input buffer */
(22) } UdsData;
```


This structure is used for all types of ports although some fields are useless for some types. The least memory consuming solution would be to arrange this structure as a union of structures, but the multiple indirections in the code to access a field in such a structure will clutter the code to much for an example.

Let's look at the fields in our structure:

- `fd` - The file descriptor of the socket associated with the port.
- `port` - The port identifier for the port which this structure corresponds to. It is needed for most `driver_XXX` calls from the driver back to the emulator.
- `lockfd` - If the socket is a listen socket, we use a separate (regular) file for two purposes:
 - We want a locking mechanism that gives no race conditions, so that we can be sure of if another Erlang node uses the listen socket name we require or if the file is only left there from a previous (crashed) session.
 - We store the *creation* serial number in the file. The *creation* is a number that should change between different instances of different Erlang emulators with the same name, so that process identifiers from one emulator won't be valid when sent to a new emulator with the same distribution name. The creation can be between 0 and 3 (two bits) and is stored in every process identifier sent to another node.
In a system with TCP based distribution, this data is kept in the *Erlang port mapper daemon* (`epmd`), which is contacted when a distributed node starts. The lock-file and a convention for the UDS listen socket's name will remove the need for `epmd` when using this distribution module. UDS is always restricted to one host, why avoiding a port mapper is easy.
- `creation` - The creation number for a listen socket, which is calculated as (the value found in the lock-file + 1) rem 4. This creation value is also written back into the lock-file, so that the next invocation of the emulator will find our value in the file.
- `type` - The current type/state of the port, which can be one of the values declared above.
- `name` - The name of the socket file (the path prefix removed), which allows for deletion (`unlink`) when the socket is closed.
- `sent` - How many bytes that have been sent over the socket. This may wrap, but that's no problem for the distribution, as the only thing that interests the Erlang distribution is if this value has changed (the Erlang `net_kernel ticker` uses this value by calling the driver to fetch it, which is done through the `driver_control` routine).
- `received` - How many bytes that are read (received) from the socket, used in similar ways as `sent`.
- `partner` - A pointer to another port structure, which is either the listen port from which this port is accepting a connection or the other way around. The "partner relation" is always bidirectional.
- `next` - Pointer to next structure in a linked list of all port structures. This list is used when accepting connections and when the driver is unloaded.
- `buffer_size`, `buffer_pos`, `header_pos`, `buffer` - data for input buffering. Refer to the source code (in the `kernel/examples` directory) for details about the input buffering. That certainly goes beyond the scope of this document.

Selected parts of the distribution driver implementation

The distribution drivers implementation is not completely covered in this text, details about buffering and other things unrelated to driver writing are not explained. Likewise are some peculiarities of the UDS protocol not explained in detail. The chosen protocol is not important.

Prototypes for the driver call-back routines can be found in the `erl_driver.h` header file.

The driver initialization routine is (usually) declared with a macro to make the driver easier to port between different operating systems (and flavours of systems). This is the only routine that has to have a well defined name. All other call-backs are reached through the driver structure. The macro to use is named `DRIVER_INIT` and takes the driver name as parameter.

```
(1) /* Beginning of linked list of ports */
(2) static UdsData *first_data;

(3) DRIVER_INIT(uds_drv)
(4) {
(5)     first_data = NULL;
(6)     return &uds_driver_entry;
(7) }
```

The routine initializes the single global data structure and returns a pointer to the driver entry. The routine will be called when `erl_dll:load_driver` is called from Erlang.

The `uds_start` routine is called when a port is opened from Erlang. In our case, we only allocate a structure and initialize it. Creating the actual socket is left to the `uds_command` routine.

```
( 1) static ErlDrvData uds_start(ErlDrvPort port, char *buff)
( 2) {
( 3)     UdsData *ud;
( 4)
( 5)     ud = ALLOC(sizeof(UdsData));
( 6)     ud->fd = -1;
( 7)     ud->lockfd = -1;
( 8)     ud->creation = 0;
( 9)     ud->port = port;
(10)     ud->type = portTypeUnknown;
(11)     ud->name = NULL;
(12)     ud->buffer_size = 0;
(13)     ud->buffer_pos = 0;
(14)     ud->header_pos = 0;
(15)     ud->buffer = NULL;
(16)     ud->sent = 0;
(17)     ud->received = 0;
(18)     ud->partner = NULL;
(19)     ud->next = first_data;
(20)     first_data = ud;
(21)
(22)     return((ErlDrvData) ud);
(23) }
```

Every data item is initialized, so that no problems will arise when a newly created port is closed (without there being any corresponding socket). This routine is called when `open_port({spawn, "uds_drv"}, [])` is called from Erlang.

The `uds_command` routine is the routine called when an Erlang process sends data to the port. All asynchronous commands when the port is in *command mode* as well as the sending of all data when the port is in *data mode* is handled in this routine. Let's have a look at it:

```
( 1) static void uds_command(ErlDrvData handle, char *buff, int buflen)
( 2) {
( 3)     UdsData *ud = (UdsData *) handle;

( 4)     if (ud->type == portTypeData || ud->type == portTypeIntermediate) {
( 5)         DEBUGF(("Passive do_send %d",buflen));
( 6)         do_send(ud, buff + 1, buflen - 1); /* XXX */
( 7)         return;
( 8)     }
( 9)     if (buflen == 0) {
(10)         return;
(11)     }
(12)     switch (*buff) {
(13)     case 'L':
(14)         if (ud->type != portTypeUnknown) {
(15)             driver_failure_posix(ud->port, ENOTSUP);
(16)             return;
(17)         }
(18)         uds_command_listen(ud,buff,buflen);
(19)         return;
(20)     case 'A':
(21)         if (ud->type != portTypeUnknown) {
(22)             driver_failure_posix(ud->port, ENOTSUP);
(23)             return;
(24)         }
(25)         uds_command_accept(ud,buff,buflen);
(26)         return;
(27)     case 'C':
(28)         if (ud->type != portTypeUnknown) {
(29)             driver_failure_posix(ud->port, ENOTSUP);
(30)             return;
(31)         }
(32)         uds_command_connect(ud,buff,buflen);
(33)         return;
(34)     case 'S':
(35)         if (ud->type != portTypeCommand) {
(36)             driver_failure_posix(ud->port, ENOTSUP);
(37)             return;
(38)         }
(39)         do_send(ud, buff + 1, buflen - 1);
(40)         return;
(41)     case 'R':
(42)         if (ud->type != portTypeCommand) {
(43)             driver_failure_posix(ud->port, ENOTSUP);
(44)             return;
(45)         }
(46)         do_rcv(ud);
(47)         return;
(48)     default:
(49)         return;
(50)     }
(51) }
```

The command routine takes three parameters; the handle returned for the port by `uds_start`, which is a pointer to the internal port structure, the data buffer and the length of the data buffer. The buffer is the data sent from Erlang (a list of bytes) converted to an C array (of bytes).

If Erlang sends i.e. the list `[$a,$b,$c]` to the port, the `bufflen` variable will be 3 and the `buff` variable will contain `{'a','b','c'}` (no null termination). Usually the first byte is used as an opcode, which is the case in our driver to (at least when the port is in command mode). The opcodes are defined as:

- 'L'<socketname>: Create and listen on socket with the given name.
- 'A'<listennumber as 32 bit bigendian>: Accept from the listen socket identified by the given identification number. The identification number is retrieved with the `uds_control` routine.
- 'C'<socketname>: Connect to the socket named <socketname>.
- 'S'<data>: Send the data <data> on the connected/accepted socket (in command mode). The sending is acked when the data has left this process.
- 'R': Receive one packet of data.

One may wonder what is meant by “one packet of data” in the 'R' command. This driver always sends data packeted with a 4 byte header containing a big endian 32 bit integer that represents the length of the data in the packet. There is no need for different packet sizes or some kind of streamed mode, as this driver is for the distribution only. One may wonder why the header word is coded explicitly in big endian when an UDS socket is local to the host. The answer simply is that I see it as a good practice when writing a distribution driver, as distribution in practice usually cross the host boundaries.

On line 4-8 we handle the case where the port is in data or intermediate mode, the rest of the routine handles the different commands. We see (first on line 15) that the routine uses the `driver_failure_posix()` routine to report errors. One important thing to remember is that the failure routines make a call to our `uds_stop` routine, which will remove the internal port data. The handle (and the casted handle `ud`) is therefore *invalid pointers* after a `driver_failure` call and we should *immediately return*. The runtime system will send exit signals to all linked processes.

The `uds_input` routine gets called when data is available on a file descriptor previously passed to the `driver_select` routine. Typically this happens when a read command is issued and no data is available. Lets look at the `do_recv` routine:

```
( 1) static void do_recv(UdsData *ud)
( 2) {
( 3)     int res;
( 4)     char *ibuf;
( 5)     for(;;) {
( 6)         if ((res = buffered_read_package(ud,&ibuf)) < 0) {
( 7)             if (res == NORMAL_READ_FAILURE) {
( 8)                 driver_select(ud->port, (ErlDrvEvent) ud->fd, DO_READ, 1);
( 9)             } else {
(10)                 driver_failure_eof(ud->port);
(11)             }
(12)             return;
(13)         }
(14)         /* Got a package */
(15)         if (ud->type == portTypeCommand) {
(16)             ibuf[-1] = 'R'; /* There is always room for a single byte
(17)                             opcode before the actual buffer
(18)                             (where the packet header was) */
(19)             driver_output(ud->port,ibuf - 1, res + 1);
(20)             driver_select(ud->port, (ErlDrvEvent) ud->fd, DO_READ,0);
```

```

(21)         return;
(22)     } else {
(23)         ibuf[-1] = DIST_MAGIC_RECV_TAG; /* XXX */
(24)         driver_output(ud->port, ibuf - 1, res + 1);
(25)         driver_select(ud->port, (ErlDrvEvent) ud->fd, DO_READ, 1);
(26)     }
(27) }
(28) }

```

The routine tries to read data until a packet is read or the `buffered_read_package` routine returns a `NORMAL_READ_FAILURE` (an internally defined constant for the module that means that the read operation resulted in an `EWOULDBLOCK`). If the port is in command mode, the reading stops when one package is read, but if it is in data mode, the reading continues until the socket buffer is empty (read failure). If no more data can be read and more is wanted (always the case when socket is in data mode) `driver_select` is called to make the `uds_input` call-back be called when more data is available for reading.

When the port is in data mode, all data is sent to Erlang in a format that suits the distribution, in fact the raw data will never reach any Erlang process, but will be translated/interpreted by the emulator itself and then delivered in the correct format to the correct processes. In the current emulator version, received data should be tagged with a single byte of 100. That's what the macro `DIST_MAGIC_RECV_TAG` is defined to. The tagging of data in the distribution will possibly change in the future.

The `uds_input` routine will handle other input events (like nonblocking `accept`), but most importantly handle data arriving at the socket by calling `do_recv`:

```

( 1) static void uds_input(ErlDrvData handle, ErlDrvEvent event)
( 2) {
( 3)     UdsData *ud = (UdsData *) handle;

( 4)     if (ud->type == portTypeListener) {
( 5)         UdsData *ad = ud->partner;
( 6)         struct sockaddr_un peer;
( 7)         int pl = sizeof(struct sockaddr_un);
( 8)         int fd;

( 9)         if ((fd = accept(ud->fd, (struct sockaddr *) &peer, &pl)) < 0) {
(10)             if (errno != EWOULDBLOCK) {
(11)                 driver_failure_posix(ud->port, errno);
(12)                 return;
(13)             }
(14)             return;
(15)         }
(16)         SET_NONBLOCKING(fd);
(17)         ad->fd = fd;
(18)         ad->partner = NULL;
(19)         ad->type = portTypeCommand;
(20)         ud->partner = NULL;
(21)         driver_select(ud->port, (ErlDrvEvent) ud->fd, DO_READ, 0);
(22)         driver_output(ad->port, "Aok", 3);
(23)         return;
(24)     }
(25)     do_recv(ud);
(26) }

```

The important line here is the last line in the function, the `do_read` routine is called to handle new input. The rest of the function handles input on a listen socket, which means that there should be possible to do an `accept` on the socket, which is also recognized as a read event.

The output mechanisms are similar to the input. Lets first look at the `do_send` routine:

```
( 1) static void do_send(UdsData *ud, char *buff, int buflen)
( 2) {
( 3)     char header[4];
( 4)     int written;
( 5)     SysIOVec iov[2];
( 6)     ErlIOVec eio;
( 7)     ErlDrvBinary *binv[] = {NULL,NULL};

( 8)     put_packet_length(header, buflen);
( 9)     iov[0].iov_base = (char *) header;
(10)     iov[0].iov_len = 4;
(11)     iov[1].iov_base = buff;
(12)     iov[1].iov_len = buflen;
(13)     eio.iov = iov;
(14)     eio.binv = binv;
(15)     eio.vsize = 2;
(16)     eio.size = buflen + 4;
(17)     written = 0;
(18)     if (driver_sizeq(ud->port) == 0) {
(19)         if ((written = writev(ud->fd, iov, 2)) == eio.size) {
(20)             ud->sent += written;
(21)             if (ud->type == portTypeCommand) {
(22)                 driver_output(ud->port, "Sok", 3);
(23)             }
(24)             return;
(25)         } else if (written < 0) {
(26)             if (errno != EWOULDBLOCK) {
(27)                 driver_failure_eof(ud->port);
(28)                 return;
(29)             } else {
(30)                 written = 0;
(31)             }
(32)         } else {
(33)             ud->sent += written;
(34)         }
(35)         /* Enqueue remaining */
(36)     }
(37)     driver_enqv(ud->port, &eio, written);
(38)     send_out_queue(ud);
(39) }
```

This driver uses the `writev` system call to send data onto the socket. A combination of `writev` and the driver output queues is very convenient. An *ErlIOVec* structure contains a *SysIOVec* (which is equivalent to the `struct iovec` structure defined in `uio.h`). The *ErlIOVec* also contains an array of *ErlDrvBinary* pointers, of the same length as the number of buffers in the I/O vector itself. One can use this to allocate the binaries for the queue “manually” in the driver, but we’ll just fill the binary array with `NULL` values (line 7), which will make the runtime system allocate it’s own buffers when we call `driver_enqv` (line 37).

The routine builds an I/O vector containing the header bytes and the buffer (the opcode has been removed and the buffer length decreased by the output routine). If the queue is empty, we'll write the data directly to the socket (or at least try to). If any data is left, it is stored in the queue and then we try to send the queue (line 38). An ack is sent when the message is delivered completely (line 22). The `send_out_queue` will send acks if the sending is completed there. If the port is in command mode, the Erlang code serializes the send operations so that only one packet can be waiting for delivery at a time. Therefore the ack can be sent simply whenever the queue is empty.

A short look at the `send_out_queue` routine:

```
( 1) static int send_out_queue(UdsData *ud)
( 2) {
( 3)     for(;;) {
( 4)         int vlen;
( 5)         SysIOVec *tmp = driver_peekq(ud->port, &vlen);
( 6)         int wrote;
( 7)         if (tmp == NULL) {
( 8)             driver_select(ud->port, (ErlDrvEvent) ud->fd, DO_WRITE, 0);
( 9)             if (ud->type == portTypeCommand) {
(10)                 driver_output(ud->port, "Sok", 3);
(11)             }
(12)             return 0;
(13)         }
(14)         if (vlen > IO_VECTOR_MAX) {
(15)             vlen = IO_VECTOR_MAX;
(16)         }
(17)         if ((wrote = writev(ud->fd, tmp, vlen)) < 0) {
(18)             if (errno == EWOULDBLOCK) {
(19)                 driver_select(ud->port, (ErlDrvEvent) ud->fd,
(20)                             DO_WRITE, 1);
(21)                 return 0;
(22)             } else {
(23)                 driver_failure_eof(ud->port);
(24)                 return -1;
(25)             }
(26)         }
(27)         driver_deq(ud->port, wrote);
(28)         ud->sent += wrote;
(29)     }
(30) }
```

What we do is simply to pick out an I/O vector from the queue (which is the whole queue as an `SysIOVec`). If the I/O vector is too long (`IO_VECTOR_MAX` is defined to 16), the vector length is decreased (line 15), otherwise the `writev` (line 17) call will fail. Writing is tried and anything written is dequeued (line 27). If the write fails with `EWOULDBLOCK` (note that all sockets are in nonblocking mode), `driver_select` is called to make the `uds_output` routine be called when there is space to write again.

We will continue trying to write until the queue is empty or the writing would block.

The routine above are called from the `uds_output` routine, which looks like this:

```
( 1) static void uds_output(ErlDrvData handle, ErlDrvEvent event)
( 2) {
( 3)     UdsData *ud = (UdsData *) handle;
```

```
( 4)   if (ud->type == portTypeConnector) {
( 5)       ud->type = portTypeCommand;
( 6)       driver_select(ud->port, (ErlDrvEvent) ud->fd, DO_WRITE, 0);
( 7)       driver_output(ud->port, "Cok",3);
( 8)       return;
( 9)   }
(10)   send_out_queue(ud);
(11) }
```

The routine is simple, it first handles the fact that the output select will concern a socket in the business of connecting (and the connecting blocked). If the socket is in a connected state it simply sends the output queue, this routine is called when there is possible to write to a socket where we have an output queue, so there is no question what to do.

The driver implements a control interface, which is a synchronous interface called when Erlang calls `erlang:driver_control/3`. This is the only interface that can control the driver when it is in data mode and it may be called with the following opcodes:

- 'C': Set port in command mode.
- 'I': Set port in intermediate mode.
- 'D': Set port in data mode.
- 'N': Get identification number for listen port, this identification number is used in an accept command to the driver, it is returned as a big endian 32 bit integer, which happens to be the file identifier for the listen socket.
- 'S': Get statistics, which is the number of bytes received, the number of bytes sent and the number of bytes pending in the output queue. This data is used when the distribution checks that a connection is alive (ticking). The statistics is returned as 3 32 bit big endian integers.
- 'T': Send a tick message, which is a packet of length 0. Ticking is done when the port is in data mode, so the command for sending data cannot be used (besides it ignores zero length packages in command mode). This is used by the ticker to send dummy data when no other traffic is present.
- 'R': Get creation number of listen socket, which is used to dig out the number stored in the lock file to differentiate between invocations of Erlang nodes with the same name.

The control interface gets a buffer to return its value in, but is free to allocate it's own buffer if the provided one is too small. Here is the code for `uds_control`:

```
( 1) static int uds_control(ErlDrvData handle, unsigned int command,
( 2)                       char* buf, int count, char** res, int res_size)
( 3) {
( 4) /* Local macro to ensure large enough buffer. */
( 5) #define ENSURE(N)                               \n( 6)      do {

(11)   UdsData *ud = (UdsData *) handle;

(12)   switch (command) {
(13)   case 'S':
(14)       {
(15)           ENSURE(13);
(16)           **res = 0;
(17)           put_packet_length((*res) + 1, ud->received);
(18)           put_packet_length((*res) + 5, ud->sent);
(19)           put_packet_length((*res) + 9, driver_sizeq(ud->port));
```



```
(20)         return 13;
(21)     }
(22) case 'C':
(23)     if (ud->type < portTypeCommand) {
(24)         return report_control_error(res, res_size, "EINVAL");
(25)     }
(26)     ud->type = portTypeCommand;
(27)     driver_select(ud->port, (ErlDrvEvent) ud->fd, DO_READ, 0);
(28)     ENSURE(1);
(29)     **res = 0;
(30)     return 1;
(31) case 'I':
(32)     if (ud->type < portTypeCommand) {
(33)         return report_control_error(res, res_size, "EINVAL");
(34)     }
(35)     ud->type = portTypeIntermediate;
(36)     driver_select(ud->port, (ErlDrvEvent) ud->fd, DO_READ, 0);
(37)     ENSURE(1);
(38)     **res = 0;
(39)     return 1;
(40) case 'D':
(41)     if (ud->type < portTypeCommand) {
(42)         return report_control_error(res, res_size, "EINVAL");
(43)     }
(44)     ud->type = portTypeData;
(45)     do_recv(ud);
(46)     ENSURE(1);
(47)     **res = 0;
(48)     return 1;
(49) case 'N':
(50)     if (ud->type != portTypeListener) {
(51)         return report_control_error(res, res_size, "EINVAL");
(52)     }
(53)     ENSURE(5);
(54)     (*res)[0] = 0;
(55)     put_packet_length((*res) + 1, ud->fd);
(56)     return 5;
(57) case 'T': /* tick */
(58)     if (ud->type != portTypeData) {
(59)         return report_control_error(res, res_size, "EINVAL");
(60)     }
(61)     do_send(ud, "", 0);
(62)     ENSURE(1);
(63)     **res = 0;
(64)     return 1;
(65) case 'R':
(66)     if (ud->type != portTypeListener) {
(67)         return report_control_error(res, res_size, "EINVAL");
(68)     }
(69)     ENSURE(2);
(70)     (*res)[0] = 0;
(71)     (*res)[1] = ud->creation;
(72)     return 2;
```

```
(73)     default:
(74)         return report_control_error(res, res_size, "EINVAL");
(75)     }
(76) #undef ENSURE
(77) }
```

The macro ENSURE (line 5 to 10) is used to ensure that the buffer is large enough for our answer. We switch on the command and take actions, there is not much to say about this routine. Worth noting is that we always has read select active on a port in data mode (achieved by calling `do_recv` on line 45), but turn off read selection in intermediate and command modes (line 27 and 36).

The rest of the driver is more or less UDS specific and not of general interest.

1.3.3 Putting it all together

To test the distribution, one can use the `net_kernel:start/1` function, which is useful as it starts the distribution on a running system, where tracing/debugging can be performed. The `net_kernel:start/1` routine takes a list as it's single argument. The lists first element should be the node name (without the "@hostname") as an atom, and the second (and last) element should be one of the atoms `shortnames` or `longnames`. In the example case `shortnames` is preferred.

For net kernel to find out which distribution module to use, the command line argument `-proto_dist` is used. The argument is followed by one or more distribution module names, with the "_dist" suffix removed, i.e. `uds_dist` as a distribution module is specified as `-proto_dist uds`.

If no `epmd` (TCP port mapper daemon) is used, one should also specify the command line option `-no_epmd`, which will make Erlang skip the `epmd` startup, both as a OS process and as an Erlang ditto.

The path to the directory where the distribution modules reside must be known at boot, which can either be achieved by specifying `-pa <path>` on the command line or by building a boot script containing the applications used for your distribution protocol (in the `uds_dist` protocol, it's only the `uds_dist` application that needs to be added to the script).

The distribution will be started at boot if all the above is specified and an `-sname <name>` flag is present at the command line, here follows two examples:

```
$ erl -pa $ERL_TOP/lib/kernel/examples/uds_dist/ebin -proto_dist uds -no_epmd
Erlang (BEAM) emulator version 5.0
```

```
Eshell V5.0 (abort with ^G)
1> net_kernel:start([bing,shortnames]).
{ok,<0.30.0>}
(bing@hador)2>
```

...

```
$ erl -pa $ERL_TOP/lib/kernel/examples/uds_dist/ebin -proto_dist uds \
-no_epmd -sname bong
Erlang (BEAM) emulator version 5.0
```

```
Eshell V5.0 (abort with ^G)
(bong@hador)1>
```

One can utilize the `ERL_FLAGS` environment variable to store the complicated parameters in:

```

$ ERL_FLAGS=-pa $ERL_TOP/lib/kernel/examples/uds_dist/ebin \
  -proto_dist uds -no_epmd
$ export ERL_FLAGS
$ erl -sname bang
Erlang (BEAM) emulator version 5.0

Eshell V5.0 (abort with ^G)
(bang@hador)1>

```

The ERL_FLAGS should preferably not include the name of the node.

1.4 The Abstract Format

This document describes the standard representation of parse trees for Erlang programs as Erlang terms. This representation is known as the *abstract format*. Functions dealing with such parse trees are `compile:forms/[1,2]` and functions in the modules `epp`, `erl_eval`, `erl_lint`, `erl_pp`, `erl_parse`, and `io`. They are also used as input and output for parse transforms (see the module `compile`).

We use the function `Rep` to denote the mapping from an Erlang source construct `C` to its abstract format representation `R`, and write $R = \text{Rep}(C)$.

The word `LINE` below represents an integer, and denotes the number of the line in the source file where the construction occurred. Several instances of `LINE` in the same construction may denote different lines.

Since operators are not terms in their own right, when operators are mentioned below, the representation of an operator should be taken to be the atom with a printname consisting of the same characters as the operator.

1.4.1 Module declarations and forms

A module declaration consists of a sequence of forms that are either function declarations or attributes.

- If `D` is a module declaration consisting of the forms `F1, ..., Fk`, then $\text{Rep}(D) = [\text{Rep}(F_1), \dots, \text{Rep}(F_k)]$.
- If `F` is an attribute `-module(Mod)`, then $\text{Rep}(F) = \{\text{attribute}, \text{LINE}, \text{module}, \text{Mod}\}$.
- If `F` is an attribute `-export([Fun1/A1, ..., Funk/Ak])`, then $\text{Rep}(F) = \{\text{attribute}, \text{LINE}, \text{export}, [\{\text{Fun}_1, \text{A}_1\}, \dots, \{\text{Fun}_k, \text{A}_k\}]\}$.
- If `F` is an attribute `-import(Mod, [Fun1/A1, ..., Funk/Ak])`, then $\text{Rep}(F) = \{\text{attribute}, \text{LINE}, \text{import}, \{\text{Mod}, [\{\text{Fun}_1, \text{A}_1\}, \dots, \{\text{Fun}_k, \text{A}_k\}]\}\}$.
- If `F` is an attribute `-compile(Options)`, then $\text{Rep}(F) = \{\text{attribute}, \text{LINE}, \text{compile}, \text{Options}\}$.
- If `F` is an attribute `-file(File, Line)`, then $\text{Rep}(F) = \{\text{attribute}, \text{LINE}, \text{file}, \{\text{File}, \text{Line}\}\}$.
- If `F` is a record declaration `-record(Name, {V1, ..., Vk})`, then $\text{Rep}(F) = \{\text{attribute}, \text{LINE}, \text{record}, \{\text{Name}, [\text{Rep}(V_1), \dots, \text{Rep}(V_k)]\}\}$. For $\text{Rep}(V)$, see below.
- If `F` is a wild attribute `-A(T)`, then $\text{Rep}(F) = \{\text{attribute}, \text{LINE}, \text{A}, \text{T}\}$.
- If `F` is a function declaration `Name Fc1 ; ... ; Name Fck`, where each `Fci` is a function clause with a pattern sequence of the same length `Arity`, then $\text{Rep}(F) = \{\text{function}, \text{LINE}, \text{Name}, \text{Arity}, [\text{Rep}(F_{c_1}), \dots, \text{Rep}(F_{c_k})]\}$.

Record fields

Each field in a record declaration may have an optional explicit default initializer expression

- If V is A , then $\text{Rep}(V) = \{\text{record_field}, \text{LINE}, \text{Rep}(A)\}$.
- If V is $A = E$, then $\text{Rep}(V) = \{\text{record_field}, \text{LINE}, \text{Rep}(A), \text{Rep}(E)\}$.

Representation of parse errors and end of file

In addition to the representations of forms, the list that represents a module declaration (as returned by functions in `erl_parse` and `epp`) may contain tuples $\{\text{error}, E\}$ and $\{\text{warning}, W\}$, denoting syntactically incorrect forms and warnings, and $\{\text{eof}, \text{LINE}\}$, denoting an end of stream encountered before a complete form had been parsed.

1.4.2 Atomic literals

There are five kinds of atomic literals, which are represented in the same way in patterns, expressions and guards:

- If L is an integer or character literal, then $\text{Rep}(L) = \{\text{integer}, \text{LINE}, L\}$.
- If L is a float literal, then $\text{Rep}(L) = \{\text{float}, \text{LINE}, L\}$.
- If L is a string literal consisting of the characters C_1, \dots, C_k , then $\text{Rep}(L) = \{\text{string}, \text{LINE}, [C_1, \dots, C_k]\}$.
- If L is an atom literal, then $\text{Rep}(L) = \{\text{atom}, \text{LINE}, L\}$.

Note that negative integer and float literals do not occur as such; they are parsed as an application of the unary negation operator.

1.4.3 Patterns

If P_s is a sequence of patterns P_1, \dots, P_k , then $\text{Rep}(P_s) = [\text{Rep}(P_1), \dots, \text{Rep}(P_k)]$. Such sequences occur as the list of arguments to a function or fun.

Individual patterns are represented as follows:

- If P is an atomic literal L , then $\text{Rep}(P) = \text{Rep}(L)$.
- If P is a compound pattern $P_1 = P_2$, then $\text{Rep}(P) = \{\text{match}, \text{LINE}, \text{Rep}(P_1), \text{Rep}(P_2)\}$.
- If P is a variable pattern V , then $\text{Rep}(P) = \{\text{var}, \text{LINE}, A\}$, where A is an atom with a printname consisting of the same characters as V .
- If P is a universal pattern $_$, then $\text{Rep}(P) = \{\text{var}, \text{LINE}, _'\}$.
- If P is a tuple pattern $\{P_1, \dots, P_k\}$, then $\text{Rep}(P) = \{\text{tuple}, \text{LINE}, [\text{Rep}(P_1), \dots, \text{Rep}(P_k)]\}$.
- If P is a nil pattern $[]$, then $\text{Rep}(P) = \{\text{nil}, \text{LINE}\}$.
- If P is a cons pattern $[P_h \mid P_t]$, then $\text{Rep}(P) = \{\text{cons}, \text{LINE}, \text{Rep}(P_h), \text{Rep}(P_t)\}$.
- If E is a binary pattern $\langle\langle P_1:\text{Size}_1/\text{TSL}_1, \dots, P_k:\text{Size}_k/\text{TSL}_k \rangle\rangle$, then $\text{Rep}(E) = \{\text{bin}, \text{LINE}, [\{\text{bin_element}, \text{LINE}, \text{Rep}(P_1), \text{Rep}(\text{Size}_1), \text{Rep}(\text{TSL}_1)\}, \dots, \{\text{bin_element}, \text{LINE}, \text{Rep}(P_k), \text{Rep}(\text{Size}_k), \text{Rep}(\text{TSL}_k)\}]\}$. For $\text{Rep}(\text{TSL})$, see below. An omitted `Size` is represented by default. An omitted `TSL` (type specifier list) is represented by default.

- If P is $P_1 \text{ Op } P_2$, where Op is a binary operator (this is either an occurrence of $++$ applied to a literal string or character list, or an occurrence of an expression that can be evaluated to a number at compile time), then $\text{Rep}(P) = \{\text{op}, \text{LINE}, \text{Op}, \text{Rep}(P_1), \text{Rep}(P_2)\}$.
- If P is $\text{Op } P_0$, where Op is a unary operator (this is an occurrence of an expression that can be evaluated to a number at compile time), then $\text{Rep}(P) = \{\text{op}, \text{LINE}, \text{Op}, \text{Rep}(P_0)\}$.
- If P is a record pattern $\#\text{Name}\{\text{Field}_1=P_1, \dots, \text{Field}_k=P_k\}$, then $\text{Rep}(P) = \{\text{record}, \text{LINE}, \text{Name}, [\{\text{record_field}, \text{LINE}, \text{Rep}(\text{Field}_1), \text{Rep}(P_1)\}, \dots, \{\text{record_field}, \text{LINE}, \text{Rep}(\text{Field}_k), \text{Rep}(P_k)\}]\}$.
- If P is $\#\text{Name}.\text{Field}$, then $\text{Rep}(P) = \{\text{record_index}, \text{LINE}, \text{Name}, \text{Rep}(\text{Field})\}$.
- If P is (P_0) , then $\text{Rep}(P) = \text{Rep}(P_0)$, i.e., patterns cannot be distinguished from their bodies.

Note that every pattern has the same source form as some expression, and is represented the same way as the corresponding expression.

1.4.4 Expressions

A body B is a sequence of expressions E_1, \dots, E_k , and $\text{Rep}(B) = [\text{Rep}(E_1), \dots, \text{Rep}(E_k)]$.

An expression E is one of the following alternatives:

- If P is an atomic literal L , then $\text{Rep}(P) = \text{Rep}(L)$.
- If E is $P = E_0$, then $\text{Rep}(E) = \{\text{match}, \text{LINE}, \text{Rep}(P), \text{Rep}(E_0)\}$.
- If E is a variable V , then $\text{Rep}(E) = \{\text{var}, \text{LINE}, A\}$, where A is an atom with a printname consisting of the same characters as V .
- If E is a tuple skeleton $\{E_1, \dots, E_k\}$, then $\text{Rep}(E) = \{\text{tuple}, \text{LINE}, [\text{Rep}(E_1), \dots, \text{Rep}(E_k)]\}$.
- If E is $[]$, then $\text{Rep}(E) = \{\text{nil}, \text{LINE}\}$.
- If E is a cons skeleton $[E_h \mid E_t]$, then $\text{Rep}(E) = \{\text{cons}, \text{LINE}, \text{Rep}(E_h), \text{Rep}(E_t)\}$.
- If E is a binary constructor $\langle\langle V_1:\text{Size}_1/\text{TSL}_1, \dots, V_k:\text{Size}_k/\text{TSL}_k \rangle\rangle$, then $\text{Rep}(E) = \{\text{bin}, \text{LINE}, [\{\text{bin_element}, \text{LINE}, \text{Rep}(V_1), \text{Rep}(\text{Size}_1), \text{Rep}(\text{TSL}_1)\}, \dots, \{\text{bin_element}, \text{LINE}, \text{Rep}(V_k), \text{Rep}(\text{Size}_k), \text{Rep}(\text{TSL}_k)\}]\}$. For $\text{Rep}(\text{TSL})$, see below. An omitted Size is represented by default. An omitted TSL (type specifier list) is represented by default.
- If E is $E_1 \text{ Op } E_2$, where Op is a binary operator, then $\text{Rep}(E) = \{\text{op}, \text{LINE}, \text{Op}, \text{Rep}(E_1), \text{Rep}(E_2)\}$.
- If E is $\text{Op } E_0$, where Op is a unary operator, then $\text{Rep}(E) = \{\text{op}, \text{LINE}, \text{Op}, \text{Rep}(E_0)\}$.
- If E is $\#\text{Name}\{\text{Field}_1=E_1, \dots, \text{Field}_k=E_k\}$, then $\text{Rep}(E) = \{\text{record}, \text{LINE}, \text{Name}, [\{\text{record_field}, \text{LINE}, \text{Rep}(\text{Field}_1), \text{Rep}(E_1)\}, \dots, \{\text{record_field}, \text{LINE}, \text{Rep}(\text{Field}_k), \text{Rep}(E_k)\}]\}$.
- If E is $E_0\#\text{Name}\{\text{Field}_1=E_1, \dots, \text{Field}_k=E_k\}$, then $\text{Rep}(E) = \{\text{record}, \text{LINE}, \text{Rep}(E_0), \text{Name}, [\{\text{record_field}, \text{LINE}, \text{Rep}(\text{Field}_1), \text{Rep}(E_1)\}, \dots, \{\text{record_field}, \text{LINE}, \text{Rep}(\text{Field}_k), \text{Rep}(E_k)\}]\}$.
- If E is $\#\text{Name}.\text{Field}$, then $\text{Rep}(E) = \{\text{record_index}, \text{LINE}, \text{Name}, \text{Rep}(\text{Field})\}$.
- If E is $E_0\#\text{Name}.\text{Field}$, then $\text{Rep}(E) = \{\text{record_field}, \text{LINE}, \text{Rep}(E_0), \text{Name}, \text{Rep}(\text{Field})\}$.
- If E is $\text{catch } E_0$, then $\text{Rep}(E) = \{\text{'catch'}, \text{LINE}, \text{Rep}(E_0)\}$.
- If E is $E_0(E_1, \dots, E_k)$, then $\text{Rep}(E) = \{\text{call}, \text{LINE}, \text{Rep}(E_0), [\text{Rep}(E_1), \dots, \text{Rep}(E_k)]\}$.

- If E is $E_m:E_0(E_1, \dots, E_k)$, then $\text{Rep}(E) = \{\text{call}, \text{LINE}, \{\text{remote}, \text{LINE}, \text{Rep}(E_m), \text{Rep}(E_0)\}, [\text{Rep}(E_1), \dots, \text{Rep}(E_k)]\}$.
- If E is a list comprehension $[E_0 \mid W_1, \dots, W_k]$, where each W_i is a generator or a filter, then $\text{Rep}(E) = \{\text{lc}, \text{LINE}, \text{Rep}(E_0), [\text{Rep}(W_1), \dots, \text{Rep}(W_k)]\}$. For $\text{Rep}(W)$, see below.
- If E is a binary comprehension $\langle\langle E_0 \mid W_1, \dots, W_k \rangle\rangle$, where each W_i is a generator or a filter, then $\text{Rep}(E) = \{\text{bc}, \text{LINE}, \text{Rep}(E_0), [\text{Rep}(W_1), \dots, \text{Rep}(W_k)]\}$. For $\text{Rep}(W)$, see below.
- If E is `begin B end`, where B is a body, then $\text{Rep}(E) = \{\text{block}, \text{LINE}, \text{Rep}(B)\}$.
- If E is `if Ic_1 ; ... ; Ic_k end`, where each Ic_i is an if clause then $\text{Rep}(E) = \{\text{'if'}, \text{LINE}, [\text{Rep}(Ic_1), \dots, \text{Rep}(Ic_k)]\}$.
- If E is `case E_0 of Cc_1 ; ... ; Cc_k end`, where E_0 is an expression and each Cc_i is a case clause then $\text{Rep}(E) = \{\text{'case'}, \text{LINE}, \text{Rep}(E_0), [\text{Rep}(Cc_1), \dots, \text{Rep}(Cc_k)]\}$.
- If E is `try B catch Tc_1 ; ... ; Tc_k end`, where B is a body and each Tc_i is a catch clause then $\text{Rep}(E) = \{\text{'try'}, \text{LINE}, \text{Rep}(B), [], [\text{Rep}(Tc_1), \dots, \text{Rep}(Tc_k)], []\}$.
- If E is `try B of Cc_1 ; ... ; Cc_k catch Tc_1 ; ... ; Tc_n end`, where B is a body, each Cc_i is a case clause and each Tc_j is a catch clause then $\text{Rep}(E) = \{\text{'try'}, \text{LINE}, \text{Rep}(B), [\text{Rep}(Cc_1), \dots, \text{Rep}(Cc_k)], [\text{Rep}(Tc_1), \dots, \text{Rep}(Tc_n)], []\}$.
- If E is `try B after A end`, where B and A are bodies then $\text{Rep}(E) = \{\text{'try'}, \text{LINE}, \text{Rep}(B), [], [], \text{Rep}(A)\}$.
- If E is `try B of Cc_1 ; ... ; Cc_k after A end`, where B and A are a bodies and each Cc_i is a case clause then $\text{Rep}(E) = \{\text{'try'}, \text{LINE}, \text{Rep}(B), [\text{Rep}(Cc_1), \dots, \text{Rep}(Cc_k)], [], \text{Rep}(A)\}$.
- If E is `try B catch Tc_1 ; ... ; Tc_k after A end`, where B and A are bodies and each Tc_i is a catch clause then $\text{Rep}(E) = \{\text{'try'}, \text{LINE}, \text{Rep}(B), [], [\text{Rep}(Tc_1), \dots, \text{Rep}(Tc_k)], \text{Rep}(A)\}$.
- If E is `try B of Cc_1 ; ... ; Cc_k catch Tc_1 ; ... ; Tc_n after A end`, where B and A are a bodies, each Cc_i is a case clause and each Tc_j is a catch clause then $\text{Rep}(E) = \{\text{'try'}, \text{LINE}, \text{Rep}(B), [\text{Rep}(Cc_1), \dots, \text{Rep}(Cc_k)], [\text{Rep}(Tc_1), \dots, \text{Rep}(Tc_n)], \text{Rep}(A)\}$.
- If E is `receive Cc_1 ; ... ; Cc_k end`, where each Cc_i is a case clause then $\text{Rep}(E) = \{\text{'receive'}, \text{LINE}, [\text{Rep}(Cc_1), \dots, \text{Rep}(Cc_k)]\}$.
- If E is `receive Cc_1 ; ... ; Cc_k after E_0 -> B_t end`, where each Cc_i is a case clause, E_0 is an expression and B_t is a body, then $\text{Rep}(E) = \{\text{'receive'}, \text{LINE}, [\text{Rep}(Cc_1), \dots, \text{Rep}(Cc_k)], \text{Rep}(E_0), \text{Rep}(B_t)\}$.
- If E is `fun Name / Arity`, then $\text{Rep}(E) = \{\text{'fun'}, \text{LINE}, \{\text{function}, \text{Name}, \text{Arity}\}\}$.
- If E is `fun Module:Name/Arity`, then $\text{Rep}(E) = \{\text{'fun'}, \text{LINE}, \{\text{function}, \text{Module}, \text{Name}, \text{Arity}\}\}$.
- If E is `fun Fc_1 ; ... ; Fc_k end` where each Fc_i is a function clause then $\text{Rep}(E) = \{\text{'fun'}, \text{LINE}, \{\text{clauses}, [\text{Rep}(Fc_1), \dots, \text{Rep}(Fc_k)]\}\}$.
- If E is `query [E_0 \mid W_1, \dots, W_k] end`, where each W_i is a generator or a filter, then $\text{Rep}(E) = \{\text{'query'}, \text{LINE}, \{\text{lc}, \text{LINE}, \text{Rep}(E_0), [\text{Rep}(W_1), \dots, \text{Rep}(W_k)]\}\}$. For $\text{Rep}(W)$, see below.
- If E is $E_0.\text{Field}$, a Mnesia record access inside a query, then $\text{Rep}(E) = \{\text{record_field}, \text{LINE}, \text{Rep}(E_0), \text{Rep}(\text{Field})\}$.
- If E is (E_0) , then $\text{Rep}(E) = \text{Rep}(E_0)$, i.e., parenthesized expressions cannot be distinguished from their bodies.

Generators and filters

When W is a generator or a filter (in the body of a list or binary comprehension), then:

- If W is a generator $P \leftarrow E$, where P is a pattern and E is an expression, then $\text{Rep}(W) = \{\text{generate}, \text{LINE}, \text{Rep}(P), \text{Rep}(E)\}$.
- If W is a generator $P \leftarrow= E$, where P is a pattern and E is an expression, then $\text{Rep}(W) = \{\text{b_generate}, \text{LINE}, \text{Rep}(P), \text{Rep}(E)\}$.
- If W is a filter E , which is an expression, then $\text{Rep}(W) = \text{Rep}(E)$.

Binary element type specifiers

A type specifier list TSL for a binary element is a sequence of type specifiers $\text{TS}_1 - \dots - \text{TS}_k$.
 $\text{Rep}(\text{TSL}) = [\text{Rep}(\text{TS}_1), \dots, \text{Rep}(\text{TS}_k)]$.

When TS is a type specifier for a binary element, then:

- If TS is an atom A , $\text{Rep}(\text{TS}) = A$.
- If TS is a couple $A : \text{Value}$ where A is an atom and Value is an integer, $\text{Rep}(\text{TS}) = \{A, \text{Value}\}$.

1.4.5 Clauses

There are function clauses, if clauses, case clauses and catch clauses.

A clause C is one of the following alternatives:

- If C is a function clause $(Ps) \rightarrow B$ where Ps is a pattern sequence and B is a body, then $\text{Rep}(C) = \{\text{clause}, \text{LINE}, \text{Rep}(Ps), [], \text{Rep}(B)\}$.
- If C is a function clause $(Ps) \text{ when } Gs \rightarrow B$ where Ps is a pattern sequence, Gs is a guard sequence and B is a body, then $\text{Rep}(C) = \{\text{clause}, \text{LINE}, \text{Rep}(Ps), \text{Rep}(Gs), \text{Rep}(B)\}$.
- If C is an if clause $Gs \rightarrow B$ where Gs is a guard sequence and B is a body, then $\text{Rep}(C) = \{\text{clause}, \text{LINE}, [], \text{Rep}(Gs), \text{Rep}(B)\}$.
- If C is a case clause $P \rightarrow B$ where P is a pattern and B is a body, then $\text{Rep}(C) = \{\text{clause}, \text{LINE}, [\text{Rep}(P)], [], \text{Rep}(B)\}$.
- If C is a case clause $P \text{ when } Gs \rightarrow B$ where P is a pattern, Gs is a guard sequence and B is a body, then $\text{Rep}(C) = \{\text{clause}, \text{LINE}, [\text{Rep}(P)], \text{Rep}(Gs), \text{Rep}(B)\}$.
- If C is a catch clause $P \rightarrow B$ where P is a pattern and B is a body, then $\text{Rep}(C) = \{\text{clause}, \text{LINE}, [\text{Rep}(\{\text{throw}, P, _ \})], [], \text{Rep}(B)\}$.
- If C is a catch clause $X : P \rightarrow B$ where X is an atomic literal or a variable pattern, P is a pattern and B is a body, then $\text{Rep}(C) = \{\text{clause}, \text{LINE}, [\text{Rep}(\{X, P, _ \})], [], \text{Rep}(B)\}$.
- If C is a catch clause $P \text{ when } Gs \rightarrow B$ where P is a pattern, Gs is a guard sequence and B is a body, then $\text{Rep}(C) = \{\text{clause}, \text{LINE}, [\text{Rep}(\{\text{throw}, P, _ \})], \text{Rep}(Gs), \text{Rep}(B)\}$.
- If C is a catch clause $X : P \text{ when } Gs \rightarrow B$ where X is an atomic literal or a variable pattern, P is a pattern, Gs is a guard sequence and B is a body, then $\text{Rep}(C) = \{\text{clause}, \text{LINE}, [\text{Rep}(\{X, P, _ \})], \text{Rep}(Gs), \text{Rep}(B)\}$.

1.4.6 Guards

A guard sequence G_s is a sequence of guards $G_{_1}; \dots; G_{_k}$, and $\text{Rep}(G_s) = [\text{Rep}(G_{_1}), \dots, \text{Rep}(G_{_k})]$. If the guard sequence is empty, $\text{Rep}(G_s) = []$.

A guard G is a nonempty sequence of guard tests G_{t_1}, \dots, G_{t_k} , and $\text{Rep}(G) = [\text{Rep}(G_{t_1}), \dots, \text{Rep}(G_{t_k})]$.

A guard test G_t is one of the following alternatives:

- If G_t is an atomic literal L , then $\text{Rep}(G_t) = \text{Rep}(L)$.
- If G_t is a variable pattern V , then $\text{Rep}(G_t) = \{\text{var}, \text{LINE}, A\}$, where A is an atom with a printname consisting of the same characters as V .
- If G_t is a tuple skeleton $\{G_{t_1}, \dots, G_{t_k}\}$, then $\text{Rep}(G_t) = \{\text{tuple}, \text{LINE}, [\text{Rep}(G_{t_1}), \dots, \text{Rep}(G_{t_k})]\}$.
- If G_t is $[\]$, then $\text{Rep}(G_t) = \{\text{nil}, \text{LINE}\}$.
- If G_t is a cons skeleton $[G_{t_h} \mid G_{t_t}]$, then $\text{Rep}(G_t) = \{\text{cons}, \text{LINE}, \text{Rep}(G_{t_h}), \text{Rep}(G_{t_t})\}$.
- If G_t is a binary constructor $\langle\langle G_{t_1}:\text{Size}_1/\text{TSL}_1, \dots, G_{t_k}:\text{Size}_k/\text{TSL}_k \rangle\rangle$, then $\text{Rep}(G_t) = \{\text{bin}, \text{LINE}, [\{\text{bin_element}, \text{LINE}, \text{Rep}(G_{t_1}), \text{Rep}(\text{Size}_1), \text{Rep}(\text{TSL}_1)\}, \dots, \{\text{bin_element}, \text{LINE}, \text{Rep}(G_{t_k}), \text{Rep}(\text{Size}_k), \text{Rep}(\text{TSL}_k)\}]\}$. For $\text{Rep}(\text{TSL})$, see above. An omitted Size is represented by `default`. An omitted TSL (type specifier list) is represented by `default`.
- If G_t is $G_{t_1} \text{ Op } G_{t_2}$, where Op is a binary operator, then $\text{Rep}(G_t) = \{\text{op}, \text{LINE}, \text{Op}, \text{Rep}(G_{t_1}), \text{Rep}(G_{t_2})\}$.
- If G_t is $\text{Op } G_{t_0}$, where Op is a unary operator, then $\text{Rep}(G_t) = \{\text{op}, \text{LINE}, \text{Op}, \text{Rep}(G_{t_0})\}$.
- If G_t is $\#\text{Name}\{\text{Field}_1=G_{t_1}, \dots, \text{Field}_k=G_{t_k}\}$, then $\text{Rep}(G_t) = \{\text{record}, \text{LINE}, \text{Name}, [\{\text{record_field}, \text{LINE}, \text{Rep}(\text{Field}_1), \text{Rep}(G_{t_1})\}, \dots, \{\text{record_field}, \text{LINE}, \text{Rep}(\text{Field}_k), \text{Rep}(G_{t_k})\}]\}$.
- If G_t is $\#\text{Name}.\text{Field}$, then $\text{Rep}(G_t) = \{\text{record_index}, \text{LINE}, \text{Name}, \text{Rep}(\text{Field})\}$.
- If G_t is $G_{t_0}\#\text{Name}.\text{Field}$, then $\text{Rep}(G_t) = \{\text{record_field}, \text{LINE}, \text{Rep}(G_{t_0}), \text{Name}, \text{Rep}(\text{Field})\}$.
- If G_t is $A(G_{t_1}, \dots, G_{t_k})$, where A is an atom, then $\text{Rep}(G_t) = \{\text{call}, \text{LINE}, \text{Rep}(A), [\text{Rep}(G_{t_1}), \dots, \text{Rep}(G_{t_k})]\}$.
- If G_t is $A_m:A(G_{t_1}, \dots, G_{t_k})$, where A_m is the atom `erlang` and A is an atom or an operator, then $\text{Rep}(G_t) = \{\text{call}, \text{LINE}, \{\text{remote}, \text{LINE}, \text{Rep}(A_m), \text{Rep}(A)\}, [\text{Rep}(G_{t_1}), \dots, \text{Rep}(G_{t_k})]\}$.
- If G_t is $\{A_m, A\}(G_{t_1}, \dots, G_{t_k})$, where A_m is the atom `erlang` and A is an atom or an operator, then $\text{Rep}(G_t) = \{\text{call}, \text{LINE}, \text{Rep}(\{A_m, A\}), [\text{Rep}(G_{t_1}), \dots, \text{Rep}(G_{t_k})]\}$.
- If G_t is (G_{t_0}) , then $\text{Rep}(G_t) = \text{Rep}(G_{t_0})$, i.e., parenthesized guard tests cannot be distinguished from their bodies.

Note that every guard test has the same source form as some expression, and is represented the same way as the corresponding expression.

1.4.7 The abstract format after preprocessing

The compilation option `debug_info` can be given to the compiler to have the abstract code stored in the `abstract_code` chunk in the BEAM file (for debugging purposes).

In OTP R9C and later, the `abstract_code` chunk will contain

```
{raw_abstract_v1, AbstractCode}
```


where `AbstractCode` is the abstract code as described in this document.

In releases of OTP prior to R9C, the abstract code after some more processing was stored in the BEAM file. The first element of the tuple would be either `abstract_v1` (R7B) or `abstract_v2` (R8B).

1.5 tty - A command line interface

`tty` is a simple command line interface program where keystrokes are collected and interpreted. Completed lines are sent to the shell for interpretation. There is a simple history mechanism, which saves previous lines. These can be edited before sending them to the shell. `tty` is started when Erlang is started with the command:

```
erl
```

`tty` operates in one of two modes:

- *normal mode*, in which lines of text can be edited and sent to the shell.
- *shell break mode*, which allows the user to kill the current shell, start multiple shells etc. Shell break mode is started by typing *Control G*.

1.5.1 Normal Mode

In normal mode keystrokes from the user are collected and interpreted by `tty`. Most of the *emacs* line editing commands are supported. The following is a complete list of the supported line editing commands.

Note: The notation `C-a` means pressing the control key and the letter `a` simultaneously. `M-f` means pressing the ESC key followed by the letter `f`.

<i>Key Sequence</i>	<i>Function</i>
C-a	Beginning of line
C-b	Backward character
M-b	Backward word
C-d	Delete character
M-d	Delete word
C-e	End of line
C-f	Forward character
M-f	Forward word
C-g	Enter shell break mode
C-k	Kill line
C-l	Redraw line
C-n	Fetch next line from the history buffer
C-p	Fetch previous line from the history buffer
C-t	Transpose characters
C-y	Insert previously killed text

Table 1.2: `tty` text editing

1.5.2 Shell Break Mode

tty enters *shell* break mode when you type *Control G*. In this mode you can:

- Kill or suspend the current shell
- Connect to a suspended shell
- Start a new shell

1.6 How to implement a driver

1.6.1 Introduction

This chapter tells you how to build your own driver for erlang.

A driver in Erlang is a library written in C, that is linked to the Erlang emulator and called from erlang. Drivers can be used when C is more suitable than Erlang, to speed things up, or to provide access to OS resources not directly accessible from Erlang.

A driver can be dynamically loaded, as a shared library (known as a DLL on windows), or statically loaded, linked with the emulator when it is compiled and linked. Only dynamically loaded drivers are described here, statically linked drivers are beyond the scope of this chapter.

When a driver is loaded it is executed in the context of the emulator, shares the same memory and the same thread. This means that all operations in the driver must be non-blocking, and that any crash in the driver will bring the whole emulator down. In short: you have to be extremely careful!

1.6.2 Sample driver

This is a simple driver for accessing a postgres database using the libpq C client library. Postgres is used because it's free and open source. For information on postgres, refer to the website www.postgres.org.

The driver is synchronous, it uses the synchronous calls of the client library. This is only for simplicity, and is generally not good, since it will halt the emulator while waiting for the database. This will be improved on below with an asynchronous sample driver.

The code is quite straight-forward: all communication between Erlang and the driver is done with `port_control/3`, and the driver returns data back using the `rbuf`.

An Erlang driver only exports one function: the driver entry function. This is defined with a macro, `DRIVER_INIT`, and returns a pointer to a C struct containing the entry points that are called from the emulator. The struct defines the entries that the emulator calls to call the driver, with a NULL pointer for entries that are not defined and used by the driver.

The `start` entry is called when the driver is opened as a port with `open_port/2`. Here we allocate memory for a user data structure. This user data will be passed every time the emulator calls us. First we store the driver handle, because it is needed in subsequent calls. We allocate memory for the connection handle that is used by LibPQ. We also set the port to return allocated driver binaries, by setting the flag `PORT_CONTROL_FLAG_BINARY`, calling `set_port_control_flags`. (This is because we don't know whether our data will fit in the result buffer of `control`, which has a default size set up by the emulator, currently 64 bytes.)

There is an entry `init` which is called when the driver is loaded, but we don't use this, since it is executed only once, and we want to have the possibility of several instances of the driver.

The `stop` entry is called when the port is closed.

The `control` entry is called from the emulator when the Erlang code calls `port_control/3`, to do the actual work. We have defined a simple set of commands: `connect` to login to the database, `disconnect` to log out and `select` to send a SQL-query and get the result. All results are returned through `rbuf`. The library `ei` in `erl_interface` is used to encode data in binary term format. The result is returned to the emulator as binary terms, so `binary_to_term` is called in Erlang to convert the result to term form.

The code is available in `pg_sync.c` in the `sample` directory of `erts`.

The driver entry contains the functions that will be called by the emulator. In our simple example, we only provide `start`, `stop` and `control`.

```
/* Driver interface declarations */
static ErlDrvData start(ErlDrvPort port, char *command);
static void stop(ErlDrvData drv_data);
static int control(ErlDrvData drv_data, unsigned int command, char *buf,
                  int len, char **rbuf, int rlen);

static ErlDrvEntry pq_driver_entry = {
    NULL,                          /* init */
    start,
    stop,
    NULL,                            /* output */
    NULL,                            /* ready_input */
    NULL,                            /* ready_output */
    "pg_sync",                       /* the name of the driver */
    NULL,                            /* finish */
    NULL,                            /* handle */
    control,
    NULL,                            /* timeout */
    NULL,                            /* outputv */
    NULL,                            /* ready_async */
    NULL,                            /* flush */
    NULL,                            /* call */
    NULL                             /* event */
};
```

We have a structure to store state needed by the driver, in this case we only need to keep the database connection.

```
typedef struct our_data_s {
    PGconn* conn;
} our_data_t;
```

These are control codes we have defined.

```
/* Keep the following definitions in alignment with the
 * defines in erl_pq_sync.erl
 */

#define DRV_CONNECT          'C'
#define DRV_DISCONNECT      'D'
#define DRV_SELECT          'S'
```

This just returns the driver structure. The macro `DRIVER_INIT` defines the only exported function. All the other functions are static, and will not be exported from the library.

```
/* INITIALIZATION AFTER LOADING */

/*
 * This is the init function called after this driver has been loaded.
 * It must not be declared static. Must return the address to
 * the driver entry.
 */

DRIVER_INIT(pq_drv)
{
    return &pq_driver_entry;
}
```

Here we do some initialization, `start` is called from `open_port`. The data will be passed to `control` and `stop`.

```
/* DRIVER INTERFACE */
static ErlDrvData start(ErlDrvPort port, char *command)
{
    our_data_t* data;

    data = (our_data_t*)driver_alloc(sizeof(our_data_t));
    data->conn = NULL;
    set_port_control_flags(port, PORT_CONTROL_FLAG_BINARY);
    return (ErlDrvData)data;
}
```

We call `disconnect` to log out from the database. (This should have been done from Erlang, but just in case.)

```
static int do_disconnect(our_data_t* data, ei_x_buff* x);

static void stop(ErlDrvData drv_data)
{
    do_disconnect((our_data_t*)drv_data, NULL);
}
```

We use the binary format only to return data to the emulator; input data is a string parameter for `connect` and `select`. The returned data consists of Erlang terms.

The functions `get_s` and `ei_x_to_new_binary` are utilities that is used to make the code shorter. `get_s` duplicates the string and zero-terminates it, since the postgres client library wants that. `ei_x_to_new_binary` takes an `ei_x_buff` buffer and allocates a binary and copies the data there. This binary is returned in `*rbuf`. (Note that this binary is freed by the emulator, not by us.)

```
static char* get_s(const char* buf, int len);
static int do_connect(const char *s, our_data_t* data, ei_x_buff* x);
static int do_select(const char* s, our_data_t* data, ei_x_buff* x);

/* Since we are operating in binary mode, the return value from control
```

```

* is irrelevant, as long as it is not negative.
*/
static int control(ErlDrvData drv_data, unsigned int command, char *buf,
                  int len, char **rbuf, int rlen)
{
    int r;
    ei_x_buff x;
    our_data_t* data = (our_data_t*)drv_data;
    char* s = get_s(buf, len);
    ei_x_new_with_version(&x);
    switch (command) {
        case DRV_CONNECT:    r = do_connect(s, data, &x); break;
        case DRV_DISCONNECT: r = do_disconnect(data, &x); break;
        case DRV_SELECT:     r = do_select(s, data, &x); break;
        default:             r = -1; break;
    }
    *rbuf = (char*)ei_x_to_new_binary(&x);
    ei_x_free(&x);
    driver_free(s);
    return r;
}

```

In `do_connect` is where we log in to the database. If the connection was successful we store the connection handle in our driver data, and return ok. Otherwise, we return the error message from postgres, and store NULL in the driver data.

```

static int do_connect(const char *s, our_data_t* data, ei_x_buff* x)
{
    PGconn* conn = PQconnectdb(s);
    if (PQstatus(conn) != CONNECTION_OK) {
        encode_error(x, conn);
        PQfinish(conn);
        conn = NULL;
    } else {
        encode_ok(x);
    }
    data->conn = conn;
    return 0;
}

```

If we are connected (if the connection handle is not NULL), we log out from the database. We need to check if a we should encode an ok, since we might get here from the `stop` function, which doesn't return data to the emulator.

```

static int do_disconnect(our_data_t* data, ei_x_buff* x)
{
    if (data->conn == NULL)
        return 0;
    PQfinish(data->conn);
    data->conn = NULL;
    if (x != NULL)
        encode_ok(x);
}

```

```
    return 0;
}
```

We execute a query and encodes the result. Encoding is done in another C module, `pg_encode.c` which is also provided as sample code.

```
static int do_select(const char* s, our_data_t* data, ei_x_buff* x)
{
    PGresult* res = PQexec(data->conn, s);
    encode_result(x, res, data->conn);
    PQclear(res);
    return 0;
}
```

Here we simply checks the result from postgres, and if it's data we encode it as lists of lists with column data. Everything from postgres is C strings, so we just use `ei_x_encode_string` to send the result as strings to Erlang. (The head of the list contains the column names.)

```
void encode_result(ei_x_buff* x, PGresult* res, PGconn* conn)
{
    int row, n_rows, col, n_cols;
    switch (PQresultStatus(res)) {
    case PGRES_TUPLES_OK:
        n_rows = PQntuples(res);
        n_cols = PQnfields(res);
        ei_x_encode_tuple_header(x, 2);
        encode_ok(x);
        ei_x_encode_list_header(x, n_rows+1);
        ei_x_encode_list_header(x, n_cols);
        for (col = 0; col < n_cols; ++col) {
            ei_x_encode_string(x, PQfname(res, col));
        }
        ei_x_encode_empty_list(x);
        for (row = 0; row < n_rows; ++row) {
            ei_x_encode_list_header(x, n_cols);
            for (col = 0; col < n_cols; ++col) {
                ei_x_encode_string(x, PQgetvalue(res, row, col));
            }
            ei_x_encode_empty_list(x);
        }
        ei_x_encode_empty_list(x);
        break;
    case PGRES_COMMAND_OK:
        ei_x_encode_tuple_header(x, 2);
        encode_ok(x);
        ei_x_encode_string(x, PQcmdTuples(res));
        break;
    default:
        encode_error(x, conn);
        break;
    }
}
```

1.6.3 Compiling and linking the sample driver

The driver should be compiled and linked to a shared library (DLL on windows). With gcc this is done with the link flags `-shared` and `-fpic`. Since we use the `ei` library we should include it too. There are several versions of `ei`, compiled for debug or non-debug and multi-threaded or single-threaded. In the makefile for the samples the `obj` directory is used for the `ei` library, meaning that we use the non-debug, single-threaded version.

1.6.4 Calling a driver as a port in Erlang

Before a driver can be called from Erlang, it must be loaded and opened. Loading is done using the `erl_ddll` module (the `erl_ddll` driver that loads dynamic driver, is actually a driver itself). If loading is ok the port can be opened with `open_port/2`. The port name must match the name of the shared library and the name in the driver entry structure.

When the port has been opened, the driver can be called. In the `pg_sync` example, we don't have any data from the port, only the return value from the `port_control`.

The following code is the Erlang part of the synchronous postgres driver, `pg_sync.erl`.

```
-module(pg_sync).

-define(DRV_CONNECT, 1).
-define(DRV_DISCONNECT, 2).
-define(DRV_SELECT, 3).

-export([connect/1, disconnect/1, select/2]).

connect(ConnectStr) ->
    case erl_ddll:load_driver(".", "pg_sync") of
        ok -> ok;
        {error, already_loaded} -> ok;
        E -> exit({error, E})
    end,
    Port = open_port({spawn, ?MODULE}, []),
    case binary_to_term(port_control(Port, ?DRV_CONNECT, ConnectStr)) of
        ok -> {ok, Port};
        Error -> Error
    end.

disconnect(Port) ->
    R = binary_to_term(port_control(Port, ?DRV_DISCONNECT, "")),
    port_close(Port),
    R.

select(Port, Query) ->
    binary_to_term(port_control(Port, ?DRV_SELECT, Query)).
```

The api is simple: `connect/1` loads the driver, opens it and logs on to the database, returning the Erlang port if successful, `select/2` sends a query to the driver, and returns the result, `disconnect/1` closes the database connection and the driver. (It does not unload it, however.) The connection string should be a connection string for postgres.

The driver is loaded with `erl_ddll:load_driver/2`, and if this is successful, or if it's already loaded, it is opened. This will call the `start` function in the driver.

We use the `port_control/3` function for all calls into the driver, the result from the driver is returned immediately, and converted to terms by calling `binary_to_term/1`. (We trust that the terms returned from the driver are well-formed, otherwise the `binary_to_term` calls could be contained in a `catch`.)

1.6.5 Sample asynchronous driver

Sometimes database queries can take long time to complete, in our `pg_sync` driver, the emulator halts while the driver is doing it's job. This is often not acceptable, since no other Erlang processes gets a chance to do anything. To improve on our postgres driver, we reimplement it using the asynchronous calls in `LibPQ`.

The asynchronous version of the driver is in the sample files `pg_async.c` and `pg_asyng.erl`.

```
/* Driver interface declarations */
static ErlDrvData start(ErlDrvPort port, char *command);
static void stop(ErlDrvData drv_data);
static int control(ErlDrvData drv_data, unsigned int command, char *buf,
                  int len, char **rbuf, int rlen);
static void ready_io(ErlDrvData drv_data, ErlDrvEvent event);

static ErlDrvEntry pq_driver_entry = {
    NULL,                /* init */
    start,
    stop,
    NULL,                /* output */
    ready_io,            /* ready_input */
    ready_io,            /* ready_output */
    "pg_async",         /* the name of the driver */
    NULL,                /* finish */
    NULL,                /* handle */
    control,
    NULL,                /* timeout */
    NULL,                /* outputv */
    NULL,                /* ready_async */
    NULL,                /* flush */
    NULL,                /* call */
    NULL,                /* event */
};

typedef struct our_data_t {
    PGconn* conn;
    ErlDrvPort port;
    int socket;
    int connecting;
} our_data_t;
```

Here some things have changed from `pg_sync.c`: we use the entry `ready_io` for `ready_input` and `ready_output` which will be called from the emulator only when there is input to be read from the socket. (Actually, the socket is used in a `select` function inside the emulator, and when the socket is signalled, indicating there is data to read, the `ready_input` entry is called. More on this below.)

Our driver data is also extended, we keep track of the socket used for communication with postgres, and also the port, which is needed when we send data to the port with `driver_output`. We have a flag `connecting` to tell whether the driver is waiting for a connection or waiting for the result of a query. (This is needed since the entry `ready_io` will be called both when connecting and when there is query result.)

```
static int do_connect(const char *s, our_data_t* data)
{
    PGconn* conn = PQconnectStart(s);
    if (PQstatus(conn) == CONNECTION_BAD) {
        ei_x_buff x;
        ei_x_new_with_version(&x);
        encode_error(&x, conn);
        PQfinish(conn);
        conn = NULL;
        driver_output(data->port, x.buff, x.index);
        ei_x_free(&x);
    }
    PQconnectPoll(conn);
    int socket = PQsocket(conn);
    data->socket = socket;
    driver_select(data->port, (ErlDrvEvent)socket, DO_READ, 1);
    driver_select(data->port, (ErlDrvEvent)socket, DO_WRITE, 1);
    data->conn = conn;
    data->connecting = 1;
    return 0;
}
```

The `connect` function looks a bit different too. We connect using the asynchronous `PQconnectStart` function. After the connection is started, we retrieve the socket for the connection with `PQsocket`. This socket is used with the `driver_select` function to wait for connection. When the socket is ready for input or for output, the `ready_io` function will be called.

Note that we only return data (with `driver_output`) if there is an error here, otherwise we wait for the connection to be completed, in which case our `ready_io` function will be called.

```
static int do_select(const char* s, our_data_t* data)
{
    data->connecting = 0;
    PGconn* conn = data->conn;
    /* if there's an error return it now */
    if (PQsendQuery(conn, s) == 0) {
        ei_x_buff x;
        ei_x_new_with_version(&x);
        encode_error(&x, conn);
        driver_output(data->port, x.buff, x.index);
        ei_x_free(&x);
    }
    /* else wait for ready_output to get results */
    return 0;
}
```

The `do_select` function initiates a select, and returns if there is no immediate error. The actual result will be returned when `ready_io` is called.

```
static void ready_io(ErlDrvData drv_data, ErlDrvEvent event)
{
    PGresult* res = NULL;
    our_data_t* data = (our_data_t*)drv_data;
    PGconn* conn = data->conn;
    ei_x_buff x;
    ei_x_new_with_version(&x);
    if (data->connecting) {
        ConnStatusType status;
        PQconnectPoll(conn);
        status = PQstatus(conn);
        if (status == CONNECTION_OK)
            encode_ok(&x);
        else if (status == CONNECTION_BAD)
            encode_error(&x, conn);
    } else {
        PQconsumeInput(conn);
        if (PQisBusy(conn))
            return;
        res = PQgetResult(conn);
        encode_result(&x, res, conn);
        PQclear(res);
        for (;;) {
            res = PQgetResult(conn);
            if (res == NULL)
                break;
            PQclear(res);
        }
    }
    if (x.index > 1) {
        driver_output(data->port, x.buff, x.index);
        if (data->connecting)
            driver_select(data->port, (ErlDrvEvent)data->socket, DO_WRITE, 0);
    }
    ei_x_free(&x);
}
```

The `ready_io` function will be called when the socket we got from postgres is ready for input or output. Here we first check if we are connecting to the database. In that case we check connection status and return ok if the connection is successful, or error if it's not. If the connection is not yet established, we simply return; `ready_io` will be called again.

If we have result from a connect, indicated that we have data in the `x` buffer, we no longer need to select on output (`ready_output`), so we remove this by calling `driver_select`.

If we're not connecting, we're waiting for results from a `PQsendQuery`, so we get the result and return it. The encoding is done with the same functions as in the earlier example.

We should add error handling here, for instance checking that the socket is still open, but this is just a simple example.

The Erlang part of the asynchronous driver consists of the sample file `pg_async.erl`.

```
-module(pg_async).
```

```

-define(DRV_CONNECT, $C).
-define(DRV_DISCONNECT, $D).
-define(DRV_SELECT, $S).

-export([connect/1, disconnect/1, select/2]).

connect(ConnectStr) ->
    case erl_ddll:load_driver(".", "pg_async") of
        ok -> ok;
        {error, already_loaded} -> ok;
        _ -> exit({error, could_not_load_driver})
    end,
    Port = open_port({spawn, ?MODULE}, [binary]),
    port_control(Port, ?DRV_CONNECT, ConnectStr),
    case return_port_data(Port) of
        ok ->
            {ok, Port};
        Error ->
            Error
    end.

disconnect(Port) ->
    port_control(Port, ?DRV_DISCONNECT, ""),
    R = return_port_data(Port),
    port_close(Port),
    R.

select(Port, Query) ->
    port_control(Port, ?DRV_SELECT, Query),
    return_port_data(Port).

return_port_data(Port) ->
    receive
        {Port, {data, Data}} ->
            binary_to_term(Data)
    end.

```

The Erlang code is slightly different, this is because we don't return the result synchronously from `port_control`, instead we get it from `driver_output` as data in the message queue. The function `return_port_data` above receives data from the port. Since the data is in binary format, we use `binary_to_term/1` to convert it to Erlang term. Note that the driver is opened in binary mode, `open_port/2` is called with the option `[binary]`. This means that data sent from the driver to the emulator is sent as binaries. Without the `binary` option, they would have been lists of integers.

1.6.6 An asynchronous driver using `driver_async`

As a final example we demonstrate the use of `driver_async`. We also use the driver term interface. The driver is written in C++. This enables us to use an algorithm from STL. We will use the `next_permutation` algorithm to get the next permutation of a list of integers. For large lists (more than 100000 elements), this will take some time, so we will perform this as an asynchronous task.

The asynchronous api for drivers are quite complicated. First of all, the work must be prepared. In our example we do this in `output`. We could have used `control` just as well, but we want some variation in

our examples. In our driver, we allocate a structure that contains all needed for the asynchronous task to do the work. This is done in the main emulator thread. Then the asynchronous function is called from a driver thread, separate from the main emulator thread. Note that the driver-functions are not reentrant, so they shouldn't be used. Finally, after the function is completed, the driver callback `ready_async` is called from the main emulator thread, this is where we return the result to Erlang. (We can't return the result from within the asynchronous function, since we can't call the driver-functions.)

The code below is from the sample file `next_perm.cc`.

The driver entry looks like before, but also contains the call-back `ready_async`.

```
static ErlDrvEntry next_perm_driver_entry = {
    NULL,                /* init */
    start,
    NULL,                /* stop */
    output,
    NULL,                /* ready_input */
    NULL,                /* ready_output */
    "next_perm",        /* the name of the driver */
    NULL,                /* finish */
    NULL,                /* handle */
    NULL,                /* control */
    NULL,                /* timeout */
    NULL,                /* outputv */
    ready_async,
    NULL,                /* flush */
    NULL,                /* call */
    NULL                 /* event */
};
```

The `output` function allocates the work-area of the asynchronous function. Since we use C++, we use a struct, and stuff the data in it. We have to copy the original data, it is not valid after we have returned from the `output` function, and the `do_perm` function will be called later, and from another thread. We return no data here, instead it will be sent later from the `ready_async` call-back.

The `async_data` will be passed to the `do_perm` function. We do not use a `async_free` function (the last argument to `driver_async`, it's only used if the task is cancelled programmatically).

```
struct our_async_data {
    bool prev;
    vector<int> data;
    our_async_data(ErlDrvPort p, int command, const char* buf, int len);
};

our_async_data::our_async_data(ErlDrvPort p, int command,
                               const char* buf, int len)
    : prev(command == 2),
      data((int*)buf, (int*)buf + len / sizeof(int))
{
}

static void do_perm(void* async_data);

static void output(ErlDrvData drv_data, char *buf, int len)
```

```

{
    if (*buf < 1 || *buf > 2) return;
    ErlDrvPort port = reinterpret_cast<ErlDrvPort>(drv_data);
    void* async_data = new our_async_data(port, *buf, buf+1, len);
    driver_async(port, NULL, do_perm, async_data, do_free);
}

```

In the `do_perm` we simply do the work, operating on the structure that was allocated in `output`.

```

static void do_perm(void* async_data)
{
    our_async_data* d = reinterpret_cast<our_async_data*>(async_data);
    if (d->prev)
        prev_permutation(d->data.begin(), d->data.end());
    else
        next_permutation(d->data.begin(), d->data.end());
}

```

In the `ready_async` function, the output is sent back to the emulator. We use the driver term format instead of `ei`. This is the only way to send Erlang terms directly to a driver, without having the Erlang code to call `binary_to_term/1`. In our simple example this works well, and we don't need to use `ei` to handle the binary term format.

When the data is returned we deallocate our data.

```

static void ready_async(ErlDrvData drv_data, ErlDrvThreadData async_data)
{
    ErlDrvPort port = reinterpret_cast<ErlDrvPort>(drv_data);
    our_async_data* d = reinterpret_cast<our_async_data*>(async_data);
    int n = d->data.size(), result_n = n*2 + 3;
    ErlDrvTermData* result = new ErlDrvTermData[result_n], * rp = result;
    for (vector<int>::iterator i = d->data.begin();
        i != d->data.end(); ++i) {
        *rp++ = ERL_DRV_INT;
        *rp++ = *i;
    }
    *rp++ = ERL_DRV_NIL;
    *rp++ = ERL_DRV_LIST;
    *rp++ = n+1;
    driver_output_term(port, result, result_n);
    delete[] result;
    delete d;
}

```

This driver is called like the others from Erlang, however, since we use `driver_output_term`, there is no need to call `binary_to_term`. The Erlang code is in the sample file `next_perm.erl`.

The input is changed into a list of integers and sent to the driver.

```

-module(next_perm).

-export([next_perm/1, prev_perm/1, load/0, all_perm/1]).

load() ->

```

```
case whereis(next_perm) of
  undefined ->
    case erl_ddll:load_driver(".", "next_perm") of
      ok -> ok;
      {error, already_loaded} -> ok;
      E -> exit(E)
    end,
    Port = open_port({spawn, "next_perm"}, []),
    register(next_perm, Port);
  _ ->
    ok
end.

list_to_integer_binaries(L) ->
  [<<I:32/integer-native>> || I <- L].

next_perm(L) ->
  next_perm(L, 1).

prev_perm(L) ->
  next_perm(L, 2).

next_perm(L, Nxt) ->
  load(),
  B = list_to_integer_binaries(L),
  port_control(next_perm, Nxt, B),
  receive
    Result ->
      Result
  end.

all_perm(L) ->
  New = prev_perm(L),
  all_perm(New, L, [New]).

all_perm(L, L, Acc) ->
  Acc;
all_perm(L, Orig, Acc) ->
  New = prev_perm(L),
  all_perm(New, Orig, [New | Acc]).
```

1.7 Inet configuration

1.7.1 Introduction

This chapter tells you how the Erlang runtime system is configured for IP communication. It also explains how you may configure it for your own particular needs by means of a configuration file. The information here is mainly intended for users with special configuration needs or problems. There should normally be no need for specific settings for Erlang to function properly on a correctly IP configured platform.

When Erlang starts up it will read the kernel variable `inetrc` which, if defined, should specify the location and name of a user configuration file. Example:

```
% erl -kernel inetrc './cfg_files/erl_inetrc'
```

Note that the usage of a `.inetrc` file, which was supported in earlier Erlang versions, is now obsolete.

A second way to specify the configuration file is to set the environment variable `ERL_INETRC` to the full name of the file. Example (bash):

```
% export ERL_INETRC=./cfg_files/erl_inetrc
```

Note that the kernel variable `inetrc` overrides this environment variable.

If no user configuration file is specified and Erlang is started in non-distributed or short name distributed mode, Erlang will use default configuration settings and a native lookup method that should work correctly under most circumstances. Erlang will not read any information from system inet configuration files (like `/etc/hosts`, `/etc/resolv.conf`, etc) in these modes.

If Erlang is started in long name distributed mode, it needs to get the domain name from somewhere and will read system inet configuration files for this information. Any hosts and resolver information found then is also recorded, but not used as long as Erlang is configured for native lookups. (The information becomes useful if the lookup method is changed to `'file'` or `'dns'`, see below).

Native lookup (system calls) is always the default resolver method. This is true for all platforms except VxWorks and OSE Delta where `'file'` or `'dns'` is used (in that order of priority).

On Windows platforms, Erlang will search the system registry rather than look for configuration files when started in long name distributed mode.

1.7.2 Configuration Data

Erlang records the following data in a local database if found in system inet configuration files (or system registry):

- Host names and addresses
- Domain name
- Nameservers
- Search domains
- Lookup method

This data may also be specified explicitly in the user configuration file. The configuration file should contain lines of configuration parameters (each terminated with a full stop). Some parameters add data to the configuration (e.g. host and nameserver), others overwrite any previous settings (e.g. domain and lookup). The user configuration file is always examined last in the configuration process, making it possible for the user to override any default values or previously made settings. Call `inet:get_rc()` to view the state of the inet configuration database.

These are the valid configuration parameters:

```
{file, Format, File}. Format = atom()
   File = string()
```

Specify a system file that Erlang should read configuration data from. `Format` tells the parser how the file should be interpreted: `resolv` (Unix `resolv.conf`), `host_conf_freebsd` (FreeBSD `host.conf`), `host_conf_bsdos` (BSDOS `host.conf`), `host_conf_linux` (Linux `host.conf`), `nsswitch_conf` (Unix `nsswitch.conf`) or `hosts` (Unix `hosts`). `File` should specify the name of the file with full path.

`{registry, Type}. Type = atom()`
Specify a system registry that Erlang should read configuration data from. Currently, `win32` is the only valid option.

`{host, IP, Aliases}. IP = tuple()`
`Aliases = [string()]`
Add host entry to the hosts table.

`{domain, Domain}. Domain = string()`
Set domain name.

`{nameserver, IP [,Port]}. IP = tuple()`
`Port = integer()`
Add address (and port, if other than default) of primary nameserver.

`{alt_nameserver, IP [,Port]}. IP = tuple()`
`Port = integer()`
Add address (and port, if other than default) of secondary nameserver.

`{search, Domains}. Domains = [string()]`
Add search domains.

`{lookup, Methods}. Methods = [atom()]`
Specify lookup methods and in which order to try them. The valid methods are: `native` (use system calls), `file` (use data retrieved from system configuration files and/or the user configuration file) or `dns` (use the Erlang DNS client for nameserver queries).

`{cache_size, Size}. Size = integer()`
Set size of resolver cache. Default is 100 DNS records.

`{cache_refresh, Time}. Time = integer()`
Set how often (in millisecond) the resolver cache is refreshed (i.e. expired DNS records are deleted). Default is 1 h.

`{timeout, Time}. Time = integer()`
Set the time to wait until retry (in millisecond) for DNS queries. Default is 2 sec.

`{retry, N}. N = integer()`
Set the number of DNS queries to try before giving up. Default is 3.

`{inet6, Bool}. Bool = true | false`
Tells the system to use IPv6. Default is false.

`{udp, Module}. Module = atom()`
Tell Erlang to use other primitive UDP module than `inet_udp`.

`{tcp, Module}. Module = atom()`
Tell Erlang to use other primitive TCP module than `inet_tcp`.

`clear_hosts`. Clear the hosts table.

`clear_ns`. Clear the list of recorded nameservers (primary and secondary).

`clear_search`. Clear the list of search domains.

1.7.3 User Configuration Example

Here follows a user configuration example.

Assume a user does not want Erlang to use the native lookup method, but wants Erlang to read all information necessary from start and use that for resolving names and addresses. In case lookup fails, Erlang should request the data from a nameserver (using the Erlang DNS client). Furthermore, DNS

records should never be cached. The user configuration file (in this example named `erl_inetrc`, stored in directory `./cfg_files`) could then look like this (Unix):

```
%% -- ERLANG INET CONFIGURATION FILE --
%% read the hosts file
{file, hosts, "/etc/hosts"}.
%% add a particular host
{host, {134,138,177,105}, ["finwe"]}.
%% read nameserver info from here
{file, resolv, "/etc/resolv.conf"}.
%% disable caching
{cache_size, 0}.
%% specify lookup method
{lookup, [file, dns]}.
```

And Erlang could, for example, be started like this:

```
% erl -sname my_node -kernel inetrc './cfg_files/erl_inetrc'
```

1.8 External Term Format

1.8.1 Introduction

The external term format is mainly used in the distribution mechanism of Erlang.

Since Erlang has a fixed number of types, there is no need for a programmer to define a specification for the external format used within some application. All Erlang terms has an external representation and the interpretation of the different terms are application specific.

In Erlang the BIF [`term_to_binary/1,2`] is used to convert a term into the external format. To convert binary data encoding a term the BIF [`binary_to_term/1`] is used.

The distribution does this implicitly when sending messages across node boundaries.

The overall format of the term format is:

1	1	N
131	Tag	Data

Table 1.3:

A compressed term looks like this:

1	1	4	N
131	80	UncompressedSize	Zlib-compressedData

Table 1.4:

Uncompressed Size (unsigned 32 bit integer in big-endian byte order) is the size of the data before it was compressed. The compressed data has the following format when it has been expanded:

1	Uncompressed Size
Tag	Data

Table 1.5:

1.8.2 SMALL_INTEGER_EXT

1	1
97	Int

Table 1.6:

Unsigned 8 bit integer.

1.8.3 INTEGER_EXT

1	4
98	Int

Table 1.7:

Signed 32 bit integer in big-endian format (i.e. MSB first)

1.8.4 FLOAT_EXT

1	31
99	Float String

Table 1.8:

A float is stored in string format. the format used in `sprintf` to format the float is `“%.20e”` (there are more bytes allocated than necessary). To unpack the float use `scanf` with format `“%lf”`.

This term is used in minor version 0 of the external format; it has been superseded by `NEW_FLOAT_EXT` [page 59].

1.8.5 ATOM_EXT

1	2	Len
100	Len	AtomName

Table 1.9:

An atom is stored with a 2 byte unsigned length in big-endian order, followed by `Len` numbers of 8 bit characters that forms the `AtomName`. Note: The maximum allowed value for `Len` is 255.

1.8.6 REFERENCE_EXT

1	N	4	1
101	Node	ID	Creation

Table 1.10:

Encode a reference object (an object generated with `make_ref/0`). The `Node` term is an encoded atom, i.e. `ATOM_EXT` [page 53], `NEW_CACHE` [page 56] or `CACHED_ATOM` [page 56]. The `ID` field contains a big-endian unsigned integer, but *should be regarded as uninterpreted data* since this field is node specific. `Creation` is a byte containing a node serial number that makes it possible to separate old (crashed) nodes from a new one.

In `ID`, only 18 bits are significant; the rest should be 0. In `Creation`, only 2 bits are significant; the rest should be 0. See `NEW_REFERENCE_EXT` [page 57].

1.8.7 PORT_EXT

1	N	4	1
102	Node	ID	Creation

Table 1.11:

Encode a port object (obtained from `open_port/2`). The `ID` is a node specific identifier for a local port. Port operations are not allowed across node boundaries. The `Creation` works just like in `REFERENCE_EXT` [page 53].

1.8.8 PID_EXT

1	N	4	4	1
103	Node	ID	Serial	Creation

Table 1.12:

Encode a process identifier object (obtained from `spawn/3` or `friends`). The `ID` and `Creation` fields works just like in `REFERENCE_EXT` [page 53], while the `Serial` field is used to improve safety. In `ID`, only 15 bits are significant; the rest should be 0.

1.8.9 SMALL_TUPLE_EXT

1	1	N
104	Arity	Elements

Table 1.13:

`SMALL_TUPLE_EXT` encodes a tuple. The `Arity` field is an unsigned byte that determines how many element that follows in the `Elements` section.

1.8.10 LARGE_TUPLE_EXT

1	4	N
105	Arity	Elements

Table 1.14:

Same as `SMALL_TUPLE_EXT` [page 54] with the exception that `Arity` is an unsigned 4 byte integer in big endian format.

1.8.11 NIL_EXT

1
106

Table 1.15:

The representation for an empty list, i.e. the Erlang syntax `[]`.

1.8.12 STRING_EXT

1	2	Len
107	Length	Characters

Table 1.16:

String does NOT have a corresponding Erlang representation, but is an optimization for sending lists of bytes (integer in the range 0-255) more efficiently over the distribution. Since the `Length` field is an unsigned 2 byte integer (big endian), implementations must make sure that lists longer than 65535 elements are encoded as `LIST_EXT` [page 55].

1.8.13 LIST_EXT

1	4		
108	Length	Elements	Tail

Table 1.17:

`Length` is the number of elements that follows in the `Elements` section. `Tail` is the final tail of the list; it is `NIL_EXT` [page 54] for a proper list, but may be anything type if the list is improper (for instance `[a|b]`).

1.8.14 BINARY_EXT

1	4	Len
109	Len	Data

Table 1.18:

Binaries are generated with bit syntax expression or with `[list_to_binary/1]`, `[term_to_binary/1]`, or as input from binary ports. The `Len` length field is an unsigned 4 byte integer (big endian).

1.8.15 SMALL_BIG_EXT

1	1	1	n
110	n	Sign	d(0) ... d(n-1)

Table 1.19:

Bignums are stored in unary form with a Sign byte that is 0 if the binum is positive and 1 if is negative. The digits are stored with the LSB byte stored first. To calculate the integer the following formula can be used:

$$B = 256$$

$$(d_0 * B^0 + d_1 * B^1 + d_2 * B^2 + \dots + d_{(N-1)} * B^{(n-1)})$$

1.8.16 LARGE_BIG_EXT

1	4	1	n
111	n	Sign	d(0) ... d(n-1)

Table 1.20:

Same as SMALL_BIG_EXT [page 55] with the difference that the length field is an unsigned 4 byte integer.

1.8.17 NEW_CACHE

1	1	2	Len
78	index	Len	Atom name

Table 1.21:

NEW_CACHE works just like ATOM_EXT [page 53], but it must also cache the atom in the atom cache in the location given by index. The atom cache is currently only used between real Erlang nodes (not between Erlang nodes and C or Java nodes).

1.8.18 CACHED_ATOM

1	1
67	index

Table 1.22:

When the atom cache is in use, index is the slot number in which the atom MUST be located.

1.8.19 NEW_REFERENCE_EXT

1	2	N	1	N'
114	Len	Node	Creation	ID ...

Table 1.23:

Node and Creation are as in REFERENCE_EXT [page 53].

ID contains a sequence of big-endian unsigned integers (4 bytes each, so N' is a multiple of 4), but should be regarded as uninterpreted data.

$N' = 4 * \text{Len}$.

In the first word (four bytes) of ID, only 18 bits are significant, the rest should be 0. In Creation, only 2 bits are significant, the rest should be 0.

NEW_REFERENCE_EXT was introduced with distribution version 4. In version 4, N' should be at most 12.

See REFERENCE_EXT [page 53]).

1.8.20 FUN_EXT

1	4	N1	N2	N3	N4	N5
117	NumFree	Pid	Module	Index	Uniq	Free vars ...

Table 1.24:

Pid is a process identifier as in PID_EXT [page 54]. It represents the process in which the fun was created.

Module is an encoded as an atom, using ATOM_EXT [page 53], NEW_CACHE [page 56] or CACHED_ATOM [page 56]. This is the module that the fun is implemented in.

Index is an integer encoded using SMALL_INTEGER_EXT [page 52] or INTEGER_EXT [page 52]. It is typically a small index into the module's fun table.

Uniq is an integer encoded using SMALL_INTEGER_EXT [page 52] or INTEGER_EXT [page 52]. Uniq is the hash value of the parse for the fun.

Free vars is NumFree number of terms, each one encoded according to its type.

1.8.21 NEW_FUN_EXT

1	4	1	16	4	4	N1	N2	N3	N4	N5
112	Size	Arity	Uniq	Index	NumFree	Module	OldIndex	OldUniq	Pid	Free Vars

Table 1.25:

This is the new encoding of internal funs: fun F/A and fun(Arg1,...) -> ... end.

`Size` is the total number of bytes, including the `Size` field.

`Arity` is the arity of the function implementing the fun.

`Uniq` is the 16 bytes MD5 of the significant parts of the Beam file.

`Index` is an index number. Each fun within a module has an unique index. `Index` is stored in big-endian byte order.

`NumFree` is the number of free variables.

`Module` is an encoded as an atom, using `ATOM_EXT` [page 53], `NEW_CACHE` [page 56] or `CACHED_ATOM` [page 56]. This is the module that the fun is implemented in.

`OldIndex` is an integer encoded using `SMALL_INTEGER_EXT` [page 52] or `INTEGER_EXT` [page 52]. It is typically a small index into the module's fun table.

`OldUniq` is an integer encoded using `SMALL_INTEGER_EXT` [page 52] or `INTEGER_EXT` [page 52]. `Uniq` is the hash value of the parse tree for the fun.

`Pid` is a process identifier as in `PID_EXT` [page 54]. It represents the process in which the fun was created.

`Free vars` is `NumFree` number of terms, each one encoded according to its type.

1.8.22 EXPORT_EXT

1	N1	N2	N3
113	Module	Function	Arity

Table 1.26:

This term is the encoding for external funs: fun M:F/A.

`Module` and `Function` are atoms (encoded using `ATOM_EXT` [page 53], `NEW_CACHE` [page 56] or `CACHED_ATOM` [page 56]).

`Arity` is an integer encoded using `SMALL_INTEGER_EXT` [page 52].

1.8.23 BIT_BINARY_EXT

1	4	1	Len
77	Len	Bits	Data

Table 1.27:

This term represents a bitstring whose length in bits is not a multiple of 8 (created using the bit syntax in R12B and later). The `Len` field is an unsigned 4 byte integer (big endian). The `Bits` field is the number of bits that are used in the last byte in the data field, counting from the most significant bit towards the least significant.

1.8.24 NEW_FLOAT_EXT

1	8
70	IEEE float

Table 1.28:

A float is stored as 8 bytes in big-endian IEEE format.

This term is used in minor version 1 of the external format.

1.9 Distribution Protocol

The description here is far from complete and will therefore be further refined in upcoming releases. The protocols both from Erlang nodes towards EPMD (Erlang Port Mapper Daemon) and between Erlang nodes are however stable and mature since many years.

The distribution protocol can be divided into four (4) parts:

- 1. Low level socket connection.
- 2. Handshake, interchange node name and authenticate.
- 3. Authentication (done by `net_kernel`).
- 4. Connected.

A node fetches the Port number of another node through the EPMD (at the other host) in order to initiate a connection request.

For each host where a distributed Erlang node is running there should also be an EPMD running. The EPMD can be started explicitly or automatically as a result of the Erlang node startup.

By default EPMD listens on port 4369.

3 and 4 are performed at the same level but the `net_kernel` disconnects the other node if it communicates using an invalid cookie (after one (1) second).

1.9.1 EPMD Protocol

The requests served by the EPMD (Erlang Port Mapper Daemon) are summarized in the figure below.




Figure 1.1: Summary of EPMD requests.

Each request *_REQ is preceded by a two-byte length field. Thus, the overall request format is:

2	n
Length	Request

Table 1.29:

Register a node in the EPMD

When a distributed node is started it registers itself in EPMD. The message `ALIVE2_REQ` described below is sent from the node towards EPMD. The response from EPMD is `ALIVE2_RESP`.

1	2	1	1	2	2	Nlen	2	Elen
120	PortNo	NodeType	Protocol	DistrvsnRange	Nlen	NodeName	Elen	Extra

Table 1.30: `ALIVE2_REQ` (120)

`PortNo` The port number on which the node accept connection requests.

`NodeType` 77 = normal Erlang node, 72 = hidden node (C-node),...

`Protocol` 0 = tcp/ip-v4, ...

`DistrvsnRange` Two bytes where MSB = Highestvsn and LSB = Lowestvsn. For erts-4.6.x (OTP-R3) the vsn = 0 For erts-4.7.x (OTP-R4) = ?????.

`Nlen` The length of the `NodeName`.

`NodeName` The `NodeName` as a string of length `Nlen`.

`Elen` The length of the `Extra` field.

`Extra` Extra field of `Elen` bytes.

The connection created to the EPMD must be kept as long as the node is a distributed node. When the connection is closed the node is automatically unregistered from the EPMD.

The response message `ALIVE2_RESP` is described below.

1	1	2
121	Result	Creation

Table 1.31: `ALIVE2_RESP` (121)

`Result` = 0 -> ok, `Result` > 0 -> error

Unregister a node from the EPMD

A node unregister itself from the EPMD by simply closing the TCP connection towards EPMD established when the node was registered.

Get the distribution port of another node

When one node wants to connect to another node it starts with a `PORT_PLEASE2_REQ` request towards EPMD on the host where the node resides in order to get the distribution port that the node listens to.

The response `PORT2_RESP` contains other valuable information such as protocol version in addition to the distribution port.

1	N
122	NodeName

Table 1.32: `PORT_PLEASE2_REQ` (122)

where $N = \text{Length} - 1$

1	1
119	Result

Table 1.33: `PORT2_RESP` (119) response indicating error, $\text{Result} > 0$.

Or

1	1	2	1	1	2	2	Nlen	2	Elen
119	Result	PortNo	NodeType	Protocol	DistrvsnRange	Nlen	NodeName	Elen	Extra

Table 1.34: `PORT2_RESP` when $\text{Result} = 0$.

If $\text{Result} > 0$, the packet only consists of [119, Result].

Get all registered names from EPMD

This request is used via the Erlang function `net_adm:names/1,2`. A TCP connection is opened towards EPMD and this request is sent.

1
110

Table 1.35: `NAMES_REQ` (110)

The response for a `NAMES_REQ` looks like this:

4	
EPMDPortNo	NodeInfo*

Table 1.36: NAMES_RESP

NodeInfo is a string written for each active node. When all NodeInfo has been written the connection is closed by EPMD.

NodeInfo is, as expressed in Erlang:

```
io:format("name ~s at port ~p~n", [NodeName, Port]).
```

Dump all data from EPMD

This request is not really used, it should be regarded as a debug feature.

1
100

Table 1.37: DUMP_REQ

The response for a DUMP_REQ looks like this:

4	
EPMDPortNo	NodeInfo*

Table 1.38: DUMP_RESP

NodeInfo is a string written for each node kept in EPMD. When all NodeInfo has been written the connection is closed by EPMD.

NodeInfo is, as expressed in Erlang:

```
io:format("active name ~s at port ~p, fd = ~p ~n",
          [NodeName, Port, Fd]).
```

or

```
io:format("old/unused name ~s at port ~p, fd = ~p~n",
          [NodeName, Port, Fd]).
```

Kill the EPMD

This request will kill the running EPMD. It is almost never used.

1
107

Table 1.39: KILL_REQ

The response fo a KILL_REQ looks like this:

2
OKString

Table 1.40: KILL_RESP

where OKString is "OK".

STOP_REQ (Not Used)

1	n
115	NodeName

Table 1.41: STOP_REQ

where $n = \text{Length} - 1$

The current implementation of Erlang does not care if the connection to the EPMD is broken.

The response for a STOP_REQ looks like this.

7
OKString

Table 1.42: STOP_RESP

where OKString is "STOPPED".

A negative response can look like this.

7
NOKString

Table 1.43: STOP_NOTOK_RESP

where NOKString is "NOEXIST".

1.9.2 Handshake

The handshake is discussed in detail in the internal documentation for the kernel (Erlang) application.

1.9.3 Protocol between connected nodes

4	1	n	m
Length	Type	ControlMsg	Message

Table 1.44:

where:

Length is equal to 1 + n + m Type is: 112 - pass through ControlMsg is a tuple passed using the external format of Erlang. Message is the message sent to another node using the '!' (in external format). But, Message is only passed in combination with a ControlMsg encoding a send ('!').

The control message is a tuple, where the first element indicates which distributed operation it encodes.

LINK {1, FromPid, ToPid}

SEND {2, Cookie, ToPid}

Note:

Message is sent as well.

EXIT {3, FromPid, ToPid, Reason}

UNLINK {4, FromPid, ToPid}

NODE_LINK {5}

REG_SEND {6, FromPid, Cookie, ToName}

Note:

Message is sent as well.

GROUP_LEADER {7, FromPid, ToPid}

EXIT2 {8, FromPid, ToPid, Reason}

1.9.4 New Ctrlmessages for distrvsn = 1 (OTP R4)

SEND_TT

{12, Cookie, ToPid, TraceToken}

Note:

Message is sent as well.

EXIT_TT

{13, FromPid, ToPid, TraceToken, Reason}

REG_SEND_TT

{16, FromPid, Cookie, ToName, TraceToken}

Note:

Message is sent as well.

EXIT2_TT

{18, FromPid, ToPid, TraceToken, Reason}

1.9.5 New Ctrlmessages for distrvsn = 2

distrvsn 2 was never used.

1.9.6 New Ctrlmessages for distrvsn = 3 (OTP R5C)

None, but the version number was increased anyway.

1.9.7 New Ctrlmessages for distrvsn = 4 (OTP R6)

These are only recognized by Erlang nodes, not by hidden nodes.

MONITOR_P

{19, FromPid, ToProc, Ref} FromPid = monitoring process ToProc = monitored process pid or name (atom)

DEMONITOR_P

{20, FromPid, ToProc, Ref} We include the FromPid just in case we want to trace this. FromPid = monitoring process ToProc = monitored process pid or name (atom)

MONITOR_P_EXIT

{21, FromProc, ToPid, Ref, Reason} FromProc = monitored process pid or name (atom) ToPid = monitoring process Reason = exit reason for the monitored process

ERTS Reference Manual

Short Summaries

- Command **epmd** [page 78] – Erlang Port Mapper Daemon
- Command **erl** [page 79] – The Erlang Emulator
- Command **erlc** [page 86] – Compiler
- Command **erlsrv** [page 89] – Run the Erlang emulator as a service on Windows NT
- Command **escript** [page 94] – Erlang scripting support
- Command **run_erl** [page 96] – Redirect Erlang input and output streams on Solaris
- Command **start** [page 98] – OTP start script example for Unix
- Command **start_erl** [page 99] – Start Erlang for embedded systems on Windows NT
- Command **werl** [page 101] – The Erlang Emulator
- C Library **erl_set_memory_block** [page 102] – Custom memory allocation for Erlang on VxWorks
- C Library **erts_alloc** [page 104] – An Erlang Run-Time System internal memory allocator library.
- Erlang Module **driver_entry** [page 110] – The driver-entry structure used by erlang drivers.
- Erlang Module **erl_driver** [page 116] – API functions for an Erlang driver

epmd

The following functions are exported:

- `epmd [-daemon]` Start a name server as a daemon
- `epmd -names` Request the names of the registered Erlang nodes on this host
- `epmd -kill` Kill the `epmdprocess`
- `epmd -help` List options

erl

The following functions are exported:

- `erl <arguments>` Start an Erlang runtime system

erlc

The following functions are exported:

- `erlc flags file1.ext file2.ext...` Compile files

erlsrv

The following functions are exported:

- `erlsrv {set | add} <service-name> [<service options>]` Add or modify an Erlang service
- `erlsrv {start | stop | disable | enable} <service-name>` Manipulate the current service status.
- `erlsrv remove <service-name>` Remove the service.
- `erlsrv list [<service-name>]` List all Erlang services or all options for one service.
- `erlsrv help` Display a brief help text

escript

The following functions are exported:

- `script-name script-arg1 script-arg2...` Run a script written in Erlang
- `escript escript-flags script-name script-arg1 script-arg2...` Run a script written in Erlang

run_erl

The following functions are exported:

- `run_erl [-daemon] pipe_dir/ log_dir "exec command [command_arguments]"` Start the Erlang emulator without attached terminal

start

The following functions are exported:

- `start [data_file]` This is an example script on how to startup the Erlang system in embedded mode on Unix.

start_erl

The following functions are exported:

- `start_erl [<erl options>] ++ [<start_erl options>]` Start the Erlang emulator with the correct release data

werl

No functions are exported.

erl_set_memory_block

The following functions are exported:

- `int erl_set_memory_block(size_t size, void *ptr, int warn_mixed_malloc, int realloc_always_moves, int use_reclaim, ...)`
Specify parameters for Erlang internal memory allocation.
- `int erl_memory_show(...)` A utility similar to VxWorks `memShow`, but for the Erlang memory area.
- `int erl_mem_info_get(MEM_PART_STATS *stats)` A utility similar to VxWorks `memPartInfoGet`, but for the Erlang memory area.

erts_alloc

No functions are exported.

driver_entry

The following functions are exported:

- `int init(void)`
[page 112] Called after loading of driver
- `int start(ErlDrvPort port, char* command)`
[page 112] Called when port is opened
- `void stop(ErlDrvData drv_data)`
[page 112] Called when port is closed
- `void output(ErlDrvData drv_data, char *buf, int len)`
[page 112] Called when port is written to
- `void ready_input(ErlDrvData drv_data, ErlDrvEvent event)`
[page 112] Called when the driver event for input or output is signaled
- `void ready_output(ErlDrvData drv_data, ErlDrvEvent event)`
[page 112] Called when the driver event for input or output is signaled
- `char *driver_name`
[page 113] The name of the driver
- `void finish(void)`
[page 113] Called just before the dynamic driver is unloaded
- `void *handle`
[page 113] Reserved, initialize to NULL
- `int control(ErlDrvData drv_data, unsigned int command, char *buf, int len, char **rbuf, int rlen)`
[page 113] Invoked with `port_control`
- `void timeout(ErlDrvData drv_data)`
[page 114] Called when timer reaches 0

- `void outputv(ErlDrvData drv_data, ErlIOVec *ev)`
[page 114] Called when the port is written to
- `void ready_async(ErlDrvData drv_data, ErlDrvThreadData thread_data)`
[page 114] Called when an asynchronous call has returned
- `int call(ErlDrvData drv_data, unsigned int command, char *buf, int len, char **rbuf, int rlen, unsigned int *flags)`
[page 114] Synchronous call with term conversion
- `void event(ErlDrvData drv_data, ErlDrvEvent event, ErlDrvEventData event_data)`
[page 114]
- `int extended_marker`
[page 114] Extended driver marker
- `int major_version`
[page 115] Major version number
- `int minor_version`
[page 115] Minor version number
- `int driver_flags`
[page 115] Driver flags
- `void *handle2`
[page 115] Reserved, initialize to NULL
- `void process_exit(ErlDrvData drv_data, ErlDrvMonitor *monitor)`
[page 115] Callback for process monitors.

erl_driver

The following functions are exported:

- `ErlDrvSysInfo`
[page 119] Information about the Erlang runtime system
- `ErlDrvBinary`
[page 120] A driver binary.
- `ErlDrvData`
[page 121] Driver specific data
- `SysIOVec`
[page 121] System I/O vector
- `ErlIOVec`
[page 121] Erlang I/O vector
- `ErlDrvMonitor`
[page 121] A monitor reference
- `ErlDrvNowData`
[page 121] A structure for holding timestamps
- `ErlDrvPDL`
[page 122] Port Data Lock
- `ErlDrvTid`
[page 122] Thread identifier
- `ErlDrvThreadOpts`
[page 122] Thread options

- ErlDrvMutex
[page 122] Mutex
- ErlDrvCond
[page 123] Condition variable
- ErlDrvRWLock
[page 123] Rwlock
- ErlDrvTSDKey
[page 123] Thread specific data key
- void driver_system_info(ErlDrvSysInfo *sys_info_ptr, size_t size)
[page 123] Get information about the Erlang runtime system
- int driver_output(ErlDrvPort port, char *buf, int len)
[page 123] Send data from driver to port owner
- int driver_output2(ErlDrvPort port, char *hbuf, int hlen, char *buf, int len)
[page 123] Send data and binary data to port owner
- int driver_output_binary(ErlDrvPort port, char *hbuf, int hlen, ErlDrvBinary* bin, int offset, int len)
[page 124] Send data from a driver binary to port owner
- int driver_outputv(ErlDrvPort port, char* hbuf, int hlen, ErlIOVec *ev, int skip)
[page 124] Send vectorized data to port owner
- int driver_vec_to_buf(ErlIOVec *ev, char *buf, int len)
[page 124] Collect data segments into a buffer
- int driver_set_timer(ErlDrvPort port, unsigned long time)
[page 124] Set a timer to call the driver
- int driver_cancel_timer(ErlDrvPort port)
[page 125] Cancel a previously set timer
- int driver_read_timer(ErlDrvPort port, unsigned long *time_left)
[page 125] Read the time left before timeout
- int driver_get_now(ErlDrvNowData *now)
[page 125] Read a system timestamp
- int driver_select(ErlDrvPort port, ErlDrvEvent event, int mode, int on)
[page 125] Provide an event for having the emulator call the driver
- void *driver_alloc(size_t size)
[page 125] Allocate memory
- void *driver_realloc(void *ptr, size_t size)
[page 126] Resize an allocated memory block
- void driver_free(void *ptr)
[page 126] Free an allocated memory block
- ErlDrvBinary* driver_alloc_binary(int size)
[page 126] Allocate a driver binary
- ErlDrvBinary* driver_realloc_binary(ErlDrvBinary *bin, int size)
[page 126] Resize a driver binary
- void driver_free_binary(ErlDrvBinary *bin)
[page 126] Free a driver binary

- `long driver_binary_get_refc(ErlDrvBinary *bin)`
[page 126] Get the reference count of a driver binary
- `long driver_binary_inc_refc(ErlDrvBinary *bin)`
[page 127] Increment the reference count of a driver binary
- `long driver_binary_dec_refc(ErlDrvBinary *bin)`
[page 127] Decrement the reference count of a driver binary
- `int driver_enq(ErlDrvPort port, char* buf, int len)`
[page 127] Enqueue data in the driver queue
- `int driver_pushq(ErlDrvPort port, char* buf, int len)`
[page 127] Push data at the head of the driver queue
- `int driver_deq(ErlDrvPort port, int size)`
[page 127] Dequeue data from the head of the driver queue
- `int driver_sizeq(ErlDrvPort port)`
[page 128] Return the size of the driver queue
- `int driver_enq_bin(ErlDrvPort port, ErlDrvBinary *bin, int offset, int len)`
[page 128] Enqueue binary in the driver queue
- `int driver_pushq_bin(ErlDrvPort port, ErlDrvBinary *bin, int offset, int len)`
[page 128] Push binary at the head of the driver queue
- `SysIOVec* driver_peekq(ErlDrvPort port, int *vlen)`
[page 128] Get the driver queue as a vector
- `int driver_enqv(ErlDrvPort port, ErlIOVec *ev, int skip)`
[page 128] Enqueue vector in the driver queue
- `int driver_pushqv(ErlDrvPort port, ErlIOVec *ev, int skip)`
[page 129] Push vector at the head of the driver queue
- `ErlDrvPDL driver_pdl_create(ErlDrvPort port)`
[page 129] Create a port data lock
- `void driver_pdl_lock(ErlDrvPDL pdl)`
[page 129] Lock port data lock
- `void driver_pdl_unlock(ErlDrvPDL pdl)`
[page 129] Unlock port data lock
- `long driver_pdl_get_refc(ErlDrvPDL pdl)`
[page 129]
- `long driver_pdl_inc_refc(ErlDrvPDL pdl)`
[page 129]
- `long driver_pdl_dec_refc(ErlDrvPDL pdl)`
[page 129]
- `int driver_monitor_process(ErlDrvPort port, ErlDrvTermData process, ErlDrvMonitor *monitor)`
[page 129] Monitor a process from a driver
- `int driver_demonitor_process(ErlDrvPort port, const ErlDrvMonitor *monitor)`
[page 130] Stop monitoring a process from a driver
- `ErlDrvTermData driver_get_monitored_process(ErlDrvPort port, const ErlDrvMonitor *monitor)`
[page 130] Retrieve the process id from a monitor

- `int driver_compare_monitors(const ErlDrvMonitor *monitor1, const ErlDrvMonitor *monitor2)`
[page 130] Compare two monitors
- `void add_driver_entry(ErlDrvEntry *de)`
[page 130] Add a driver entry
- `int remove_driver_entry(ErlDrvEntry *de)`
[page 130] Remove a driver entry
- `char* erl_errno_id(int error)`
[page 131] Get erlang error atom name from error number
- `void set_busy_port(ErlDrvPort port, int on)`
[page 131] Signal or unsignal port as busy
- `void set_port_control_flags(ErlDrvPort port, int flags)`
[page 131] Set flags on how to handle control entry function
- `int driver_failure_eof(ErlDrvPort port)`
[page 131] Fail with EOF
- `int driver_failure_atom(ErlDrvPort port, char *string)`
[page 131] Fail with error
- `int driver_failure_posix(ErlDrvPort port, int error)`
[page 131] Fail with error
- `int driver_failure(ErlDrvPort port, int error)`
[page 131] Fail with error
- `ErlDrvTermData driver_connected(ErlDrvPort port)`
[page 131] Return the port owner process
- `ErlDrvTermData driver_caller(ErlDrvPort port)`
[page 131] Return the process making the driver call
- `int driver_output_term(ErlDrvPort port, ErlDrvTermData* term, int n)`
[page 132] Send term data from driver to port owner
- `ErlDrvTermData driver_mk_atom(char* string)`
[page 134] Make an atom from a name
- `ErlDrvTermData driver_mk_port(ErlDrvPort port)`
[page 134] Make a erlang term port from a port
- `int driver_send_term(ErlDrvPort port, ErlDrvTermData receiver, ErlDrvTermData* term, int n)`
[page 134] Send term data to other process than port owner process
- `long driver_async (ErlDrvPort port, unsigned int* key, void (*async_invoke)(void*), void* async_data, void (*async_free)(void*))`
[page 134] Perform an asynchronous call within a driver
- `int driver_async_cancel(long id)`
[page 135] Cancel an asynchronous call
- `int driver_lock_driver(ErlDrvPort port)`
[page 135] Make sure the driver is never unloaded
- `ErlDrvPort driver_create_port(ErlDrvPort port, ErlDrvTermData owner_pid, char* name, ErlDrvData drv_data)`
[page 135] Create a new port (driver instance)
- `int erl_drv_thread_create(char *name, ErlDrvTid *tid, void * (*func)(void *), void *arg, ErlDrvThreadOpts *opts)`
[page 136] Create a thread

- `ErlDrvThreadOpts *erl_drv_thread_opts_create(char *name)`
[page 137] Create thread options
- `void erl_drv_thread_opts_destroy(ErlDrvThreadOpts *opts)`
[page 137] Destroy thread options
- `void erl_drv_thread_exit(void *exit_value)`
[page 137] Terminate calling thread
- `int erl_drv_thread_join(ErlDrvTid tid, void **exit_value)`
[page 137] Join with another thread
- `ErlDrvTid erl_drv_thread_self(void)`
[page 138] Get the thread identifier of the current thread
- `int erl_drv_equal_tids(ErlDrvTid tid1, ErlDrvTid tid2)`
[page 138] Compare thread identifiers for equality
- `ErlDrvMutex *erl_drv_mutex_create(char *name)`
[page 138] Create a mutex
- `void erl_drv_mutex_destroy(ErlDrvMutex *mtx)`
[page 138] Destroy a mutex
- `void erl_drv_mutex_lock(ErlDrvMutex *mtx)`
[page 139] Lock a mutex
- `int erl_drv_mutex_trylock(ErlDrvMutex *mtx)`
[page 139] Try lock a mutex
- `void erl_drv_mutex_unlock(ErlDrvMutex *mtx)`
[page 139] Unlock a mutex
- `ErlDrvCond *erl_drv_cond_create(char *name)`
[page 139] Create a condition variable
- `void erl_drv_cond_destroy(ErlDrvCond *cnd)`
[page 140] Destroy a condition variable
- `void erl_drv_cond_signal(ErlDrvCond *cnd)`
[page 140] Signal on a condition variable
- `void erl_drv_cond_broadcast(ErlDrvCond *cnd)`
[page 140] Broadcast on a condition variable
- `void erl_drv_cond_wait(ErlDrvCond *cnd, ErlDrvMutex *mtx)`
[page 140] Wait on a condition variable
- `ErlDrvRWLock *erl_drv_rwlock_create(char *name)`
[page 141] Create an rwlock
- `void erl_drv_rwlock_destroy(ErlDrvRWLock *rwlck)`
[page 141] Destroy an rwlock
- `void erl_drv_rwlock_rlock(ErlDrvRWLock *rwlck)`
[page 141] Read lock an rwlock
- `int erl_drv_rwlock_tryrlock(ErlDrvRWLock *rwlck)`
[page 142] Try to read lock an rwlock
- `void erl_drv_rwlock_runlock(ErlDrvRWLock *rwlck)`
[page 142] Read unlock an rwlock
- `void erl_drv_rwlock_rwlock(ErlDrvRWLock *rwlck)`
[page 142] Read/Write lock an rwlock
- `int erl_drv_rwlock_tryrwlock(ErlDrvRWLock *rwlck)`
[page 143] Try to read/write lock an rwlock

- `void erl_drv_rwlock_rwlock(ErlDrvRWLock *rwlock)`
[page 143] Read/Write unlock an rwlock
- `int erl_drv_tsd_key_create(char *name, ErlDrvTSDKey *key)`
[page 143] Create a thread specific data key
- `void erl_drv_tsd_key_destroy(ErlDrvTSDKey key)`
[page 144] Destroy a thread specific data key
- `void erl_drv_tsd_set(ErlDrvTSDKey key, void *data)`
[page 144] Set thread specific data
- `void *erl_drv_tsd_get(ErlDrvTSDKey key)`
[page 144] Get thread specific data
- `int erl_drv_putenv(char *key, char *value)`
[page 144] Set the value of an environment variable
- `int erl_drv_getenv(char *key, char *value, size_t *value_size)`
[page 145] Get the value of an environment variable

epmd

Command

This daemon acts as a name server on all hosts involved in distributed Erlang computations. When an Erlang node starts, the node has a name and it obtains an address from the host OS kernel. The name and the address are sent to the `epmd` daemon running on the local host. In a TCP/IP environment, the address consists of the IP address and a port number. The name of the node is an atom on the form of `Name@Node`. The job of the `epmd` daemon is to keep track of which node name listens on which address. Hence, `epmd` map symbolic node names to machine addresses.

The daemon is started automatically by the Erlang start-up script.

The program `epmd` can also be used for a variety of other purposes, for example checking the DNS (Domain Name System) configuration of a host.

Exports

`epmd [-daemon]`

Starts a name server as a daemon. If it has no argument, the `epmd` runs as a normal program with the controlling terminal of the shell in which it is started. Normally, it should run as a daemon.

`epmd -names`

Requests the names of the local Erlang nodes `epmd` has registered.

`epmd -kill`

Kills the `epmd` process.

`epmd -help`

Write short info about the usage including some debugging options not listed here.

Logging

On some operating systems *syslog* will be used for error reporting when `epmd` runs as an daemon. To enable the error logging you have to edit `/etc/syslog.conf` file and add an entry

```
!epmd
*.*<TABs>/var/log/epmd.log
```

where `<TABs>` are at least one real tab character. Spaces will silently be ignored.

erl

Command

The `erl` program starts an Erlang runtime system. The exact details (for example, whether `erl` is a script or a program and which other programs it calls) are system-dependent.

Windows users probably wants to use the `werl` program instead, which runs in its own window with scrollbars and supports command-line editing. The `erl` program on Windows provides no line editing in its shell, and on Windows 95 there is no way to scroll back to text which has scrolled off the screen. The `erl` program must be used, however, in pipelines or if you want to redirect standard input or output.

Exports

`erl <arguments>`

Starts an Erlang runtime system.

The arguments can be divided into *emulator flags*, *flags* and *plain arguments*:

- Any argument starting with the character `+` is interpreted as an emulator flag [page 83].
As indicated by the name, emulator flags controls the behavior of the emulator.
- Any argument starting with the character `-` (hyphen) is interpreted as a flag [page 80] which should be passed to the Erlang part of the runtime system, more specifically to the `init` system process, see [init(3)].
The `init` process itself interprets some of these flags, the *init flags*. It also stores any remaining flags, the *user flags*. The latter can be retrieved by calling `init:get_argument/1`.
It can be noted that there are a small number of `-` flags which now actually are emulator flags, see the description below.
- Plain arguments are not interpreted in any way. They are also stored by the `init` process and can be retrieved by calling `init:get_plain_arguments/0`. Plain arguments can occur before the first flag, or after a `--` flag. Additionally, the flag `-extra` causes everything that follows to become plain arguments.

Example:

```
% erl +W w -sname arnie +R 9 -s my_init -extra +bertie
(arnie@host)1> init:get_argument(sname).
{ok, [{"arnie"}]}
(arnie@host)2> init:get_plain_arguments().
["+bertie"]
```

Here `+W w` and `+R 9` are emulator flags. `-s my_init` is an init flag, interpreted by `init`. `-sname arnie` is a user flag, stored by `init`. It is read by Kernel and will cause the Erlang runtime system to become distributed. Finally, everything after `-extra` (that is, `+bertie`) is considered as plain arguments.

```
% erl -myflag 1
1> init:get_argument(myflag).
{ok, [{"1"]}}
2> init:get_plain_arguments().
[]
```

Here the user flag `-myflag 1` is passed to and stored by the `init` process. It is a user defined flag, presumably used by some user defined application.

Flags

In the following list, init flags are marked (init flag). Unless otherwise specified, all other flags are user flags, for which the values can be retrieved by calling `init:get_argument/1`. Note that the list of user flags is not exhaustive, there may be additional, application specific flags which instead are documented in the corresponding application documentation.

- (init flag)** Everything following `--` up to the next flag (`-flag` or `+flag`) is considered plain arguments and can be retrieved using `init:get_plain_arguments/0`.
- `-Application Par Val` Sets the application configuration parameter `Par` to the value `Val` for the application `Application`, see `[app(4)]` and `[application(3)]`.
- `-args_file FileName` Command line arguments are read from the file `FileName`. The arguments read from the file replace the `'-args_file FileName'` flag on the resulting command line.
 The file `FileName` should be a plain text file and may contain comments and command line arguments. A comment begins with a `#` character and continues until next end of line character. Backslash (`\`) is used as quoting character. All command line arguments accepted by `erl` are allowed, also the `-args_file FileName` flag. Be careful not to cause circular dependencies between files containing the `-args_file` flag, though.
 The `-extra` flag is treated specially. Its scope ends at the end of the file. Arguments following an `-extra` flag are moved on the command line into the `-extra` section, i.e. the end of the command line following after an `-extra` flag.
- `-async_shell_start` The initial Erlang shell does not read user input until the system boot procedure has been completed (Erlang 5.4 and later). This flag disables the start synchronization feature and lets the shell start in parallel with the rest of the system.
- `-boot File` Specifies the name of the boot file, `File.boot`, which is used to start the system. See `[init(3)]`. Unless `File` contains an absolute path, the system searches for `File.boot` in the current and `$ROOT/bin` directories.
 Defaults to `$ROOT/bin/start.boot`.
- `-boot_var Var Dir` If the boot script contains a path variable `Var` other than `$ROOT`, this variable is expanded to `Dir`. Used when applications are installed in another directory than `$ROOT/lib`, see `[systools:make_script/1,2]`.
- `-code_path_cache` Enables the code path cache of the code server, see `[code(3)]`.

- `compile Mod1 Mod2 ...` Compiles the specified modules and then terminates (with non-zero exit code if the compilation of some file did not succeed). Implies `-noinput`. Not recommended - use `erlc` [page 86] instead.
- `config Config` Specifies the name of a configuration file, `Config.config`, which is used to configure applications. See [app(4)] and [application(3)].
- `connect_all false` If this flag is present, `global` will not maintain a fully connected network of distributed Erlang nodes, and then `global` name registration cannot be used. See [global(3)].
- `cookie Cookie` Obsolete flag without any effect and common misspelling for `-setcookie`. Use `-setcookie` instead.
- `detached` Starts the Erlang runtime system detached from the system console. Useful for running daemons and background processes.
- `emu_args` Useful for debugging. Prints out the actual arguments sent to the emulator.
- `env Variable Value` Sets the host OS environment variable `Variable` to the value `Value` for the Erlang runtime system. Example:

```
% erl -env DISPLAY gin:0
```

In this example, an Erlang runtime system is started with the `DISPLAY` environment variable set to `gin:0`.

- `eval Expr` (**init flag**) Makes `init` evaluate the expression `Expr`, see [init(3)].
- `extra` (**init flag**) Everything following `-extra` is considered plain arguments and can be retrieved using `init:get_plain_arguments/0`.
- `heart` Starts heart beat monitoring of the Erlang runtime system. See [heart(3)].
- `hidden` Starts the Erlang runtime system as a hidden node, if it is run as a distributed node. Hidden nodes always establish hidden connections to all other nodes except for nodes in the same global group. Hidden connections are not published on neither of the connected nodes, i.e. neither of the connected nodes are part of the result from `nodes/0` on the other node. See also hidden global groups, [global_group(3)].
- `hosts Hosts` Specifies the IP addresses for the hosts on which Erlang boot servers are running, see [erl_boot_server(3)]. This flag is mandatory if the `-loader inet` flag is present.
The IP addresses must be given in the standard form (four decimal numbers separated by periods, for example "150.236.20.74". Hosts names are not acceptable, but a broadcast address (preferably limited to the local network) is.
- `id Id` Specifies the identity of the Erlang runtime system. If it is run as a distributed node, `Id` must be identical to the name supplied together with the `-sname` or `-name` flag.
- `init_debug` Makes `init` write some debug information while interpreting the boot script.
- `instr` (**emulator flag**) Selects an instrumented Erlang runtime system (virtual machine) to run, instead of the ordinary one. When running an instrumented runtime system, some resource usage data can be obtained and analysed using the module `instrument`. Functionally, it behaves exactly like an ordinary Erlang runtime system.

- loader *Loader* Specifies the method used by `erl_prim_loader` to load Erlang modules into the system. See `[erl_prim_loader(3)]`. Two *Loader* methods are supported, `efile` and `inet`. `efile` means use the local file system, this is the default. `inet` means use a boot server on another machine, and the `-id`, `-hosts` and `-setcookie` flags must be specified as well. If *Loader* is something else, the user supplied *Loader* port program is started.
- make Makes the Erlang runtime system invoke `make:all()` in the current working directory and then terminate. See `[make(3)]`. Implies `-noinput`.
- man *Module* Displays the manual page for the Erlang module *Module*. Only supported on Unix.
- mode `interactive` | `embedded` Indicates if the system should load code dynamically (`interactive`), or if all code should be loaded during system initialization (`embedded`), see `[code(3)]`. Defaults to `interactive`.
- name *Name* Makes the Erlang runtime system into a distributed node. This flag invokes all network servers necessary for a node to become distributed. See `[net_kernel(3)]`. It is also ensured that `epmd` runs on the current host before Erlang is started. See `epmd(1)` [page 78].
The name of the node will be `Name@Host`, where *Host* is the fully qualified host name of the current host. For short names, use the `-sname` flag instead.
- noinput Ensures that the Erlang runtime system never tries to read any input. Implies `-noshell`.
- noshell Starts an Erlang runtime system with no shell. This flag makes it possible to have the Erlang runtime system as a component in a series of UNIX pipes.
- nostick Disables the sticky directory facility of the Erlang code server, see `[code(3)]`.
- oldshell Invokes the old Erlang shell from Erlang 3.3. The old shell can still be used.
- pa *Dir1 Dir2 ...* Adds the specified directories to the beginning of the code path, similar to `code:add_pathsa/1`. See `[code(3)]`. As an alternative to `-pa`, if several directories are to be prepended to the code and the directories have a common parent directory, that parent directory could be specified in the `ERL_LIBS` environment variable. See `[code(3)]`.
- pz *Dir1 Dir2 ...* Adds the specified directories to the end of the code path, similar to `code:add_pathsz/1`. See `[code(3)]`.
- remsh *Node* Starts Erlang with a remote shell connected to *Node*.
- rsh *Program* Specifies an alternative to `rsh` for starting a slave node on a remote host. See `[slave(3)]`.
- run *Mod* [*Func* [*Arg1, Arg2, ...*]] (**init flag**) Makes `init` call the specified function. *Func* defaults to `start`. If no arguments are provided, the function is assumed to be of arity 0. Otherwise it is assumed to be of arity 1, taking the list [*Arg1, Arg2, ...*] as argument. All arguments are passed as strings. See `[init(3)]`.
- s *Mod* [*Func* [*Arg1, Arg2, ...*]] (**init flag**) Makes `init` call the specified function. *Func* defaults to `start`. If no arguments are provided, the function is assumed to be of arity 0. Otherwise it is assumed to be of arity 1, taking the list [*Arg1, Arg2, ...*] as argument. All arguments are passed as atoms. See `[init(3)]`.
- setcookie *Cookie* Sets the magic cookie of the node to *Cookie*, see `[erlang:set_cookie/2]`.
- shutdown_time *Time* Specifies how long time (in milliseconds) the `init` process is allowed to spend shutting down the system. If *Time* ms have elapsed, all processes still existing are killed. Defaults to `infinity`.

- `-sname Name` Makes the Erlang runtime system into a distributed node, similar to `-name`, but the host name portion of the node name `Name@Host` will be the short name, not fully qualified.
This is sometimes the only way to run distributed Erlang if the DNS (Domain Name System) is not running. There can be no communication between nodes running with the `-sname` flag and those running with the `-name` flag, as node names must be unique in distributed Erlang systems.
- `-smp [enable|auto|disable]` `-smp enable` and `-smp` starts the Erlang runtime system with SMP support enabled. This may fail if no runtime system with SMP support is available. `-smp auto` starts the Erlang runtime system with SMP support enabled if it is available and more than one logical processor are detected. `-smp disable` starts a runtime system without SMP support. By default `-smp auto` will be used unless a conflicting parameter has been passed, then `-smp disable` will be used. Currently only the `-hybrid` parameter conflicts with `-smp auto`.
NOTE: The runtime system with SMP support will not be available on all supported platforms. See also the `+S` [page 84] flag.
- `-version(emulator flag)` Makes the emulator print out its version number. The same as `erl +V`.

Emulator Flags

`erl` invokes the code for the Erlang emulator (virtual machine), which supports the following flags:

- `+a size` Suggested stack size, in kilowords, for threads in the `async-thread` pool. Valid range is 16-8192 kilowords. The default suggested stack size is 16 kilowords, i.e. 64 kilobyte on 32-bit architectures. This small default size has been chosen since the amount of `async-threads` might be quite large. The default size is enough for drivers delivered with Erlang/OTP, but might not be sufficiently large for other dynamically linked in drivers that use the `driver_async()` [page 134] functionality. Note that the value passed is only a suggestion, and it might even be ignored on some platforms.
- `+A size` Sets the number of threads in `async thread` pool, valid range is 0-1024. Default is 0.
- `+B [c | d | i]` The `c` option makes `Ctrl-C` interrupt the current shell instead of invoking the emulator break handler. The `d` option (same as specifying `+B` without an extra option) disables the break handler. The `i` option makes the emulator ignore any break signal.
If the `c` option is used with `oldshell` on Unix, `Ctrl-C` will restart the shell process rather than interrupt it.
Note that on Windows, this flag is only applicable for `werl`, not `erl (oldshell)`.
Note also that `Ctrl-Break` is used instead of `Ctrl-C` on Windows.
- `+c` Disable compensation for sudden changes of system time.
Normally, `erlang:now/0` will not immediately reflect sudden changes in the system time, in order to keep timers (including `receive-after`) working. Instead, the time maintained by `erlang:now/0` is slowly adjusted towards the new system time. (Slowly means in one percent adjustments; if the time is off by one minute, the time will be adjusted in 100 minutes.)
When the `+c` option is given, this slow adjustment will not take place. Instead `erlang:now/0` will always reflect the current system time. Note that timers are

based on `erlang:now/0`. If the system time jumps, timers then time out at the wrong time.

- +h `Size` Sets the default heap size of processes to the size `Size`.
- +K `true | false` Enables or disables the kernel poll functionality if the emulator supports it. Default is `false` (disabled). If the emulator does not support kernel poll, and the +K flag is passed to the emulator, a warning is issued at startup.
- +l Enables auto load tracing, displaying info while loading code.
- +MFlag `Value` Memory allocator specific flags, see `erts_alloc(3)` [page 104] for further information.
- +P `Number` Sets the maximum number of concurrent processes for this system. `Number` must be in the range 16..134217727. Default is 32768.
- +R `ReleaseNumber` Sets the compatibility mode.
 The distribution mechanism is not backwards compatible by default. This flag sets the emulator in compatibility mode with an earlier Erlang/OTP release `ReleaseNumber`. The release number must be in the range 7..<current release>. This limits the emulator, making it possible for it to communicate with Erlang nodes (as well as C- and Java nodes) running that earlier release.
 For example, an R10 node is not automatically compatible with an R9 node, but R10 nodes started with the +R 9 flag can co-exist with R9 nodes in the same distributed Erlang system, they are R9-compatible.
 Note: Make sure all nodes (Erlang-, C-, and Java nodes) of a distributed Erlang system is of the same Erlang/OTP release, or from two different Erlang/OTP releases X and Y, where *all* Y nodes have compatibility mode X.
 For example: A distributed Erlang system can consist of R10 nodes, or of R9 nodes and R9-compatible R10 nodes, but not of R9 nodes, R9-compatible R10 nodes and “regular” R10 nodes, as R9 and “regular” R10 nodes are not compatible.
- +r Force ets memory block to be moved on realloc.
- +S `Number` Sets the number of scheduler threads to use when SMP support has been enabled. Valid range is 1-1024. If the Erlang runtime system is able to determine the number of processor cores available, the default value will equal the this value; otherwise, the default value will be one.
 This flag will be ignored if the emulator doesn't have SMP support enabled (see the `-smp` [page 83] flag).
- +T `Level` Enables modified timing and sets the modified timing level. Currently valid range is 0-9. The timing of the runtime system will change. A high level usually means a greater change than a low level. Changing the timing can be very useful for finding timing related bugs.
 Currently, modified timing affects the following:
 - Process spawning** A process calling `spawn`, `spawn_link`, `spawn_monitor`, or `spawn_opt` will be scheduled out immediately after completing the call. When higher modified timing levels are used, the caller will also sleep for a while after being scheduled out.
 - Context reductions** The amount of reductions a process is allowed to use before being scheduled out is increased or reduced.
 - Input reductions** The amount of reductions performed before checking I/O is increased or reduced.

NOTE: Performance will suffer when modified timing is enabled. This flag is *only* intended for testing and debugging. Also note that `return_to` and `return_from` trace messages will be lost when tracing on the spawn BIFs. This flag may be removed or changed at any time without prior notice.

- +v Makes the emulator print out its version number.
- +v Verbose.
- +W w | i Sets the mapping of warning messages for `error_logger`. Messages sent to the error logger using one of the warning routines can be mapped either to errors (default), warnings (+W w), or info reports (+W i). The current mapping can be retrieved using `error_logger:warning_map/0`. See `[error_logger(3)]` for further information.

Environment variables

- `ERL_CRASH_DUMP` If the emulator needs to write a crash dump, the value of this variable will be the file name of the crash dump file. If the variable is not set, the name of the crash dump file will be `erl_crash.dump` in the current directory.
- `ERL_CRASH_DUMP_NICE` *Unix systems:* If the emulator needs to write a crash dump, it will use the value of this variable to set the nice value for the process, thus lowering its priority. The allowable range is 1 through 39 (higher values will be replaced with 39). The highest value, 39, will give the process the lowest priority.
- `ERL_CRASH_DUMP_SECONDS` *Unix systems:* This variable gives the number of seconds that the emulator will be allowed to spend writing a crash dump. When the given number of seconds have elapsed, the emulator will be terminated by a SIGALRM signal.
- `ERL_AFLAGS` The content of this environment variable will be added to the beginning of the command line for `erl`.
The `-extra` flag is treated specially. Its scope ends at the end of the environment variable content. Arguments following an `-extra` flag are moved on the command line into the `-extra` section, i.e. the end of the command line following after an `-extra` flag.
- `ERL_ZFLAGS` and `ERL_FLAGS` The content of these environment variables will be added to the end of the command line for `erl`.
The `-extra` flag is treated specially. Its scope ends at the end of the environment variable content. Arguments following an `-extra` flag are moved on the command line into the `-extra` section, i.e. the end of the command line following after an `-extra` flag.
- `ERL_LIBS` This environment variable contains a list of additional library directories that the code server will search for applications and add to the code path. See `[code(3)]`.

SEE ALSO

`[init(3)]`, `[erl_prim_loader(3)]`, `[erl_boot_server(3)]`, `[code(3)]`, `[application(3)]`, `[heart(3)]`, `[net_kernel(3)]`, `[auth(3)]`, `[make(3)]`, `epmd(1)` [page 78], `erts.alloc(3)` [page 104]

erlc

Command

The `erlc` program provides a common way to run all compilers in the Erlang system. Depending on the extension of each input file, `erlc` will invoke the appropriate compiler. Regardless of which compiler is used, the same flags are used to provide parameters such as include paths and output directory.

Exports

```
erlc flags file1.ext file2.ext...
```

`Er1c` compiles one or more files. The files must include the extension, for example `.er1` for Erlang source code, or `.yr1` for Yecc source code. `Er1c` uses the extension to invoke the correct compiler.

Generally Useful Flags

The following flags are supported:

- I** *directory* Instructs the compiler to search for include files in the specified directory. When encountering an `-include` or `-include_dir` directive, the compiler searches for header files in the following directories:
 1. ".", the current working directory of the file server;
 2. the base name of the compiled file;
 3. the directories specified using the `-I` option. The directory specified last is searched first.
- o** *directory* The directory where the compiler should place the output files. If not specified, output files will be placed in the current working directory.
- Dname** Defines a macro.
- Dname=value** Defines a macro with the given value. The value can be any Erlang term. Depending on the platform, the value may need to be quoted if the shell itself interprets certain characters. On Unix, terms which contain tuples and list must be quoted. Terms which contain spaces must be quoted on all platforms.
- Wnumber** Sets warning level to *number*. Default is 1. Use `-W0` to turn off warnings.
- W** Same as `-W1`. Default.
- v** Enables verbose output.
- b** *output-type* Specifies the type of output file. Generally, *output-type* is the same as the file extension of the output file but without the period. This option will be ignored by compilers that have a single output format.

- hybrid** Compile using the hybrid-heap emulator. This is mainly useful for compiling native code, which needs to be compiled with the same run-time system that it should be run on.
- smp** Compile using the SMP emulator. This is mainly useful for compiling native code, which needs to be compiled with the same run-time system that it should be run on.
- Signals that no more options will follow. The rest of the arguments will be treated as file names, even if they start with hyphens.
- +**term** A flag starting with a plus ('+') rather than a hyphen will be converted to an Erlang term and passed unchanged to the compiler. For instance, the `export_all` option for the Erlang compiler can be specified as follows:

```
erlc +export_all file.erl
```

Depending on the platform, the value may need to be quoted if the shell itself interprets certain characters. On Unix, terms which contain tuples and list must be quoted. Terms which contain spaces must be quoted on all platforms.

Special Flags

The flags in this section are useful in special situations such as re-building the OTP system.

- pa *directory*** Appends *directory* to the front of the code path in the invoked Erlang emulator. This can be used to invoke another compiler than the default one.
- pz *directory*** Appends *directory* to the code path in the invoked Erlang emulator.

Supported Compilers

- .erl** Erlang source code. It generates a `.beam` file.
 - The options `-P`, `-E`, and `-S` are equivalent to `+'P'`, `+'E'`, and `+'S'`, except that it is not necessary to include the single quotes to protect them from the shell.
 - Supported options: `-I`, `-o`, `-D`, `-v`, `-W`, `-b`.
- .yrl** Yecc source code. It generates an `.erl` file.
 - Use the `-I` option with the name of a file to use that file as a customized prologue file (the `includefile` option).
 - Supported options: `-o`, `-v`, `-I`, `-W` (see above).
- .mib** MIB for SNMP. It generates a `.bin` file.
 - Supported options: `-I`, `-o`, `-W`.
- .bin** A compiled MIB for SNMP. It generates a `.hrl` file.
 - Supported options: `-o`, `-v`.
- .rel** Script file. It generates a boot file.
 - Use the `-I` to name directories to be searched for application files (equivalent to the path in the option list for `systools:make_script/2`).
 - Supported options: `-o`.

.asn1 ASN1 file.

Creates an `.er1`, `.hr1`, and `.asn1db` file from an `.asn1` file. Also compiles the `.er1` using the Erlang compiler unless the `+noobj` options is given.

Supported options: `-I`, `-o`, `-b`, `-W`.

.idl IC file.

Runs the IDL compiler.

Supported options: `-I`, `-o`.

Environment Variables

ERLC_EMULATOR The command for starting the emulator. Default is `erl` in the same directory as the `erlc` program itself, or if it doesn't exist, `erl` in any of the directories given in the `PATH` environment variable.

SEE ALSO

`erl(1)` [page 79], `[compile(3)]`, `[yecc(3)]`, `[snmp(3)]`

erlsrv

Command

This utility is specific to Windows NT/2000/XP. It allows Erlang emulators to run as services on the Windows system, allowing embedded systems to start without any user needing to log in. The emulator started in this way can be manipulated through the Windows services applet in a manner similar to other services.

As well as being the actual service, `erlsrv` also provides a command line interface for registering, changing, starting and stopping services.

To manipulate services, the logged in user should have Administrator privileges on the machine. The Erlang machine itself is (default) run as the local administrator. This can be changed with the Services applet in Windows.

The processes created by the service can, as opposed to normal services, be “killed” with the task manager. Killing an emulator that is started by a service will trigger the “OnFail” action specified for that service, which may be a reboot.

The following parameters may be specified for each Erlang service:

- **StopAction:** This tells `erlsrv` how to stop the Erlang emulator. Default is to kill it (Win32 `TerminateProcess`), but this action can specify any Erlang shell command that will be executed in the emulator to make it stop. The emulator is expected to stop within 30 seconds after the command is issued in the shell. If the emulator is not stopped, it will report a running state to the service manager.
- **OnFail:** This can be either of `reboot`, `restart`, `restart_always` or `ignore` (the default). In case of `reboot`, the NT system is rebooted whenever the emulator stops (a more simple form of watchdog), this could be useful for less critical systems, otherwise use the heart functionality to accomplish this. The `restart` value makes the Erlang emulator be restarted (with whatever parameters are registered for the service at the occasion) when it stops. If the emulator stops again within 10 seconds, it is not restarted to avoid an infinite loop which could completely hang the NT system. `restart_always` is similar to `restart`, but does not try to detect cyclic restarts, it is expected that some other mechanism is present to avoid the problem. The default (`ignore`) just reports the service as stopped to the service manager whenever it fails, it has to be manually restarted.
On a system where release handling is used, this should always be set to `ignore`. Use `heart` to restart the service on failure instead.
- **Machine:** The location of the Erlang emulator. The default is the `erl.exe` located in the same directory as `erlsrv.exe`. Do not specify `werl.exe` as this emulator, it will not work.
If the system uses release handling, this should be set to a program similar to `start_erl.exe`.

- **Env:** Specifies an *additional* environment for the emulator. The environment variables specified here are added to the system wide environment block that is normally present when a service starts up. Variables present in both the system wide environment and in the service environment specification will be set to the value specified in the service.
- **WorkDir:** The working directory for the Erlang emulator, has to be on a local drive (there are no network drives mounted when a service starts). Default working directory for services is %SystemDrive%%SystemPath%. Debug log files will be placed in this directory.
- **Priority:** The process priority of the emulator, this can be one of `realtime`, `high`, `low` or `default` (the default). Real-time priority is not recommended, the machine will possibly be inaccessible to interactive users. High priority could be used if two Erlang nodes should reside on one dedicated system and one should have precedence over the other. Low process priority may be used if interactive performance should not be affected by the emulator process.
- **SName or Name:** Specifies the short or long node-name of the Erlang emulator. The Erlang services are always distributed, default is to use the service name as (short) node-name.
- **DebugType:** Can be one of `none` (default), `new`, `reuse` or `console`. Specifies that output from the Erlang shell should be sent to a “debug log”. The log file is named `<servicename>.debug` or `<servicename>.debug.<N>`, where `<N>` is an integer between 1 and 99. The log-file is placed in the working directory of the service (as specified in `WorkDir`). The `reuse` option always reuses the same log file (`<servicename>.debug`) and the `new` option uses a separate log file for every invocation of the service (`<servicename>.debug.<N>`). The `console` option opens an interactive Windows console window for the Erlang shell of the service. The `console` option automatically disables the `StopAction` and a service started with an interactive console window will not survive logouts, `OnFail` actions do not work with debug-consoles either. If no `DebugType` is specified (`none`), the output of the Erlang shell is discarded.
The `consoleDebugType` is *not in any way* intended for production. It is *only* a convenient way to debug Erlang services during development. The `new` and `reuse` options might seem convenient to have in a production system, but one has to take into account that the logs will grow indefinitely during the systems lifetime and there is no way, short of restarting the service, to truncate those logs. In short, the `DebugType` is intended for debugging only. Logs during production are better produced with the standard Erlang logging facilities.
- **Args:** Additional arguments passed to the emulator startup program `erl.exe` (or `start_erl.exe`). Arguments that cannot be specified here are `-noinput` (`StopActions` would not work), `-name` and `-sname` (they are specified in any way). The most common use is for specifying cookies and flags to be passed to `init:boot()` (`-s`).

The naming of the service in a system that uses release handling has to follow the convention `NodeName_Release`, where `NodeName` is the first part of the Erlang nodename (up to, but not including the “@”) and `Release` is the current release of the application.

Exports

```
erlsrv {set | add} <service-name> [<service options>]
```

The set and add commands adds or modifies a Erlang service respectively. The simplest form of an add command would be completely without options in which case all default values (described above) apply. The service name is mandatory.

Every option can be given without parameters, in which case the default value is applied. Values to the options are supplied *only* when the default should not be used (i.e. `erlsrv set myservice -prio -arg` sets the default priority and removes all arguments).

The following service options are currently available:

- stopaction** [<erlang shell command>]] Defines the StopAction, the command given to the Erlang shell when the service is stopped. Default is none.
- onfail** [{reboot | restart | restart_always}]] Specifies the action to take when the Erlang emulator stops unexpectedly. Default is to ignore.
- machine** [<erl-command>]] The complete path to the Erlang emulator, never use the `werl` program for this. Default is the `erl.exe` in the same directory as `erlsrv.exe`. When release handling is used, this should be set to a program similar to `start_erl.exe`.
- env** [<variable>[=<value>]] ...] Edits the environment block for the service. Every environment variable specified will add to the system environment block. If a variable specified here has the same name as a system wide environment variable, the specified value overrides the system wide. Environment variables are added to this list by specifying `<variable>=<value>` and deleted from the list by specifying `<variable>` alone. The environment block is automatically sorted. Any number of `-env` options can be specified in one command. Default is to use the system environment block unmodified (except for two additions, see below [page 92]).
- workdir** [<directory>]] The initial working directory of the Erlang emulator. Default is the system directory.
- priority** [{low | high | realtime}]] The priority of the Erlang emulator. The default is the Windows default priority.
- {-**sname** | -n[ame]} [<node-name>]] The node-name of the Erlang machine, distribution is mandatory. Default is `-sname <service name>`.
- debugtype** [{new | reuse | console}]] Specifies where shell output should be sent, default is that shell output is discarded. To be used only for debugging.
- args** [<limited erl arguments>]] Additional arguments to the Erlang emulator, avoid `-noinput`, `-noshell` and `-sname/-name`. Default is no additional arguments. Remember that the services cookie file is not necessarily the same as the interactive users. The service runs as the local administrator. All arguments should be given together in one string, use double quotes (") to give an argument string containing spaces and use quoted quotes (\") to give an quote within the argument string if necessary.

```
erlsrv {start | stop | disable | enable} <service-name>
```

These commands are only added for convenience, the normal way to manipulate the state of a service is through the control panels services applet. The `start` and `stop` commands communicates with the service manager for stopping and starting a service. The commands wait until the service is actually stopped or started. When disabling a service, it is not stopped, the disabled state will not take effect until the service actually is stopped. Enabling a service sets it in automatic mode, that is started at boot. This command cannot set the service to manual.

```
erlsrv remove <service-name>
```

This command removes the service completely with all its registered options. It will be stopped before it is removed.

```
erlsrv list [<service-name>]
```

If no service name is supplied, a brief listing of all Erlang services is presented. If a service-name is supplied, all options for that service are presented.

```
erlsrv help
```

ENVIRONMENT

The environment of an Erlang machine started as a service will contain two special variables, `ERLSRV_SERVICE_NAME`, which is the name of the service that started the machine and `ERLSRV_EXECUTABLE` which is the full path to the `erlsrv.exe` that can be used to manipulate the service. This will come in handy when defining a heart command for your service. A command file for restarting a service will simply look like this:

```
@echo off
%ERLSRV_EXECUTABLE% stop %ERLSRV_SERVICE_NAME%
%ERLSRV_EXECUTABLE% start %ERLSRV_SERVICE_NAME%
```

This command file is then set as heart command.

The environment variables can also be used to detect that we are running as a service and make port programs react correctly to the control events generated on logout (see below).

PORT PROGRAMS

When a program runs in the service context, it has to handle the control events that is sent to every program in the system when the interactive user logs off. This is done in different ways for programs running in the console subsystem and programs running as window applications. An application which runs in the console subsystem (normal for port programs) uses the win32 function `SetConsoleCtrlHandler` to a control handler that returns `TRUE` in answer to the `CTRL_LOGOFF_EVENT`. Other applications just forward `WM_ENDSESSION` and `WM_QUERYENDSESSION` to the default window procedure. Here is a brief example in C of how to set the console control handler:

```
#include <windows.h>
/*
** A Console control handler that ignores the log off events,
** and lets the default handler take care of other events.
*/
BOOL WINAPI service_aware_handler(DWORD ctrl){
    if(ctrl == CTRL_LOGOFF_EVENT)
        return TRUE;
    return FALSE;
}

void initialize_handler(void){
    char buffer[2];
    /*
    * We assume we are running as a service if this
    * environment variable is defined
    */
    if(GetEnvironmentVariable("ERLSRV_SERVICE_NAME",buffer,
                              (DWORD) 2)){
        /*
        ** Actually set the control handler
        */
        SetConsoleCtrlHandler(&service_aware_handler, TRUE);
    }
}
```

NOTES

Even though the options are described in a Unix-like format, the case of the options or commands is not relevant, and the “/” character for options can be used as well as the “-” character.

Note that the program resides in the emulators bin-directory, not in the bin-directory directly under the Erlang root. The reasons for this are the subtle problem of upgrading the emulator on a running system, where a new version of the runtime system should not need to overwrite existing (and probably used) executables.

To easily manipulate the Erlang services, put the `<erlang_root>\erts-<version>\bin` directory in the path instead of `<erlang_root>\bin`. The `erlsrv` program can be found from inside Erlang by using the `os:find_executable/1` Erlang function.

For release handling to work, use `start_erl` as the Erlang machine. It is also worth mentioning again that the name of the service is significant (see above [page 90]).

SEE ALSO

`start_erl(1)`, `release_handler(3)`

escript

Command

escript provides support for running short Erlang programs without having to compile them first and an easy way to retrieve the command line arguments.

Exports

```
script-name script-arg1 script-arg2...
```

```
escript escript-flags script-name script-arg1 script-arg2...
```

escript runs a script written in Erlang.

Here follows an example.

```
$ cat factorial
#!/usr/bin/env escript
%% -*- erlang -*-
main([String]) ->
    try
        N = list_to_integer(String),
        F = fac(N),
        io:format("factorial ~w = ~w\n", [N,F])
    catch
        _:_ ->
            usage()
    end;
main(_) ->
    usage().

usage() ->
    io:format("usage: factorial integer\n"),
    halt(1).

fac(0) -> 1;
fac(N) -> N * fac(N-1).
$ factorial 5
factorial 5 = 120
$ factorial
usage: factorial integer
$ factorial five
usage: factorial integer
```

Note that there should not be any module declaration in an Erlang script file. Instead, the first line is usually the interpreter line which invokes escript. If you invoke escript like this

```
$ escript factorial 5
```

the contents of the first line does not matter, but it cannot contain Erlang code as it will be *ignored*.

If you know the location of the `escript` executable, the first line can directly give the path to `escript`. For instance:

```
#!/usr/local/bin/escript
```

As any other kind of scripts, Erlang scripts will not work on Unix platforms if they execution bit for the script file is not set. (Use `chmod +x script-name` to turn on the execution bit.)

An Erlang script file must always contain the function `main/1`. When the script is run, the `main/1` will be called with a list of strings representing the arguments given to the script (not changed or interpreted in any way).

Call `escript:script_name/0` from your script to retrieve the pathname of the script (the pathname is usually, but not always, absolute).

It is not necessary to export the `main/1` function.

By default, the script will be interpreted. You can force it to be compiled by including the following line somewhere in the script file:

```
-mode(compile).
```

Pre-processor directives in the script files are ignored, with the exception for the `-include_lib` directive. For instance, use

```
-include_lib("kernel/include/file.hrl").
```

to include the record definitions for the records used by the `file:read_file_info/1` function.

Pre-defined macros (such as `?MODULE`) will not work. A script does not have module name, so BIFs such as `[spawn/3]` that require a module name cannot be used. Instead, use a BIF that take a fun, such as `[spawn/1]`.

The script will be checked for syntactic and semantic correctness before being run. If there are warnings (such as unused variables), they will be printed and the script will still be run. If there are errors, they will be printed and the script will not be run and its exit status will be 127.

If the `main/1` function in the script returns successfully, the exit status for the script will be 0. If an exception is generated during execution, a short message will be printed and the script terminated with exit status 127.

To return your own non-zero exit code, call `halt(ExitCode)`; for instance:

```
halt(1).
```

Options accepted by escript

- s Only perform a syntactic and semantic check of the script file. Warnings and errors (if any) are written to the standard output, but the script will not be run. The exit status will be 0 if there were no errors, and 127 otherwise.

run_erl

Command

This describes the `run_erl` program specific to Solaris/Linux. This program redirect the standard input and standard output streams so that all output can be logged. It also let the program `to_erl` connect to the Erlang console making it possible to monitor and debug an embedded system remotely.

You can read more about the use in the *Embedded System User's Guide*.

Exports

```
run_erl [-daemon] pipe_dir/ log_dir "exec command [command_arguments]"
```

The `run_erl` program arguments are:

-daemon This option is highly recommended. It makes `run_erl` run in the background completely detached from any controlling terminal and the command returns to the caller immediately. Without this option, `run_erl` must be started using several tricks in the shell to detach it completely from the terminal in use when starting it. The option must be the first argument to `run_erl` on the command line.

pipe_dir This is where to put the named pipe, usually `/tmp/`. It shall be suffixed by a `/` (slash), i.e. not `/tmp/epipes`, but `/tmp/epipes/`.

log_dir This is where the log files are written. There will be one log file, `run_erl.log` that log progress and warnings from the `run_erl` program itself and there will be up to five log files at maximum 100KB each (both number of logs and sizes can be changed by environment variables, see below) with the content of the standard streams from and to the command. When the logs are full `run_erl` will delete and reuse the oldest log file.

"exec command [command_arguments]" In the third argument `command` is the to execute where everything written to `stdin` and `stdout` is logged to `log_dir`.

Notes concerning the log files

While running, `run_erl` (as stated earlier) sends all output, uninterpreted, to a log file. The file is called `erlang.log.N`, where `N` is a number. When the log is "full", default after 100KB, `run_erl` starts to log in file `erlang.log.(N+1)`, until `N` reaches a certain number (default 5), where after `N` starts at 1 again and the oldest files start getting overwritten. If no output comes from the erlang shell, but the erlang machine still seems to be alive, an "ALIVE" message is written to the log, it is a timestamp and is written, by default, after 15 minutes of inactivity. Also, if output from erlang is logged but it's been more than 5 minutes (default) since last time we got anything from erlang, a timestamp is written in the log. The "ALIVE" messages look like this:

```
==== ALIVE <date-time-string>
```

while the other timestamps look like this:

```
==== <date-time-string>
```

The `date-time-string` is the date and time the message is written, default in local time (can be changed to GMT if one wants to) and is formatted with the ANSI-C function `strftime` using the format string `%a %b %e %T %Z %Y`, which produces messages on the line of `==== ALIVE Thu May 15 10:13:36 MEST 2003`, this can be changed, see below.

Environment variables

The following environment variables are recognized by `run_ertl` and change the logging behavior. Also see the notes above to get more info on how the log behaves.

RUN_ERTL_LOG_ALIVE_MINUTES How long to wait for output (in minutes) before writing an "ALIVE" message to the log. Default is 15, can never be less than 1.

RUN_ERTL_LOG_ACTIVITY_MINUTES How long erlang need to be inactive before output will be preceded with a timestamp. Default is `RUN_ERTL_LOG_ALIVE_MINUTES div 3`, but never less than 1.

RUN_ERTL_LOG_ALIVE_FORMAT Specifies another format string to be used in the `strftime` C library call. i.e specifying this to `"%e-%b-%Y, %T %Z"` will give log messages with timestamps looking like `15-May-2003, 10:23:04 MET` etc. See the documentation for the C library function `strftime` for more information. Default is `"%a %b %e %T %Z %Y"`.

RUN_ERTL_LOG_ALIVE_IN_UTC If set to anything else than "0", it will make all times displayed by `run_ertl` to be in UTC (GMT,CET,MET, without DST), rather than in local time. This does not affect data coming from erlang, only the logs output directly by `run_ertl`. The application `sasl` can be modified accordingly by setting the erlang application variable `utc_log` to `true`.

RUN_ERTL_LOG_GENERATIONS Controls the number of log files written before older files are being reused. Default is 5, minimum is 2, maximum is 1000.

RUN_ERTL_LOG_MAXSIZE The size (in bytes) of a log file before switching to a new log file. Default is 100000, minimum is 1000 and maximum is approximately 2^{30} .

SEE ALSO

`start(1)`, `start_ertl(1)`

start

Command

This describes the `start` script that is an example script on how to startup the Erlang system in embedded mode on Unix.

You can read more about the use in the *Embedded System User's Guide*.

Exports

```
start [ data_file ]
```

In the example there is one argument

data_file Optional, specifies what `start_erl.data` file to use.

There is also an environment variable `RELDIR` that can be set prior to calling this example that set the directory where to find the release files.

SEE ALSO

`run_erl(1)`, `start_erl(1)`

start_erl

Command

This describes the `start_erl` program specific to Windows NT. Although there exists programs with the same name on other platforms, their functionality is not the same.

The `start_erl` program is distributed both in compiled form (under `<Erlang root>\erts-<version>\bin`) and in source form (under `<Erlang root>\erts-<version>\src`). The purpose of the source code is to make it possible to easily customize the program for local needs, such as cyclic restart detection etc. There is also a “make”-file, written for the `nmake` program distributed with Microsoft Visual C++. The program can however be compiled with any Win32 C compiler (possibly with slight modifications).

The purpose of the program is to aid release handling on Windows NT. The program should be called by the `erlsrv` program, read up the release data file `start_erl.data` and start Erlang. Certain options to `start_erl` are added and removed by the release handler during upgrade with emulator restart (more specifically the `-data` option).

Exports

```
start_erl [<erl options>] ++ [<start_erl options>]
```

The `start_erl` program in its original form recognizes the following options:

- `++` Mandatory, delimits `start_erl` options from normal Erlang options. Everything on the command line *before* the `++` is interpreted as options to be sent to the `erl` program. Everything *after* `++` is interpreted as options to `start_erl` itself.
- `-reldir <release root>` Mandatory if the environment variable `RELDIR` is not specified. Tells `start_erl` where the root of the release tree is placed in the file-system (like `<Erlang root>\releases`). The `start_erl.data` file is expected to be placed in this directory (if not otherwise specified).
- `-data <data file name>` Optional, specifies another data file than `start_erl.data` in the `<release root>`. It is specified relative to the `<release root>` or absolute (including drive letter etc.). This option is used by the release handler during upgrade and should not be used during normal operation. The release data file should not normally be named differently.
- `-bootflags <boot flags file name>` Optional, specifies a file name relative to actual release directory (that is the subdirectory of `<release root>` where the `.boot` file etc. are placed). The contents of this file is appended to the command line when Erlang is started. This makes it easy to start the emulator with different options for different releases.

NOTES

As the source code is distributed, it can easily be modified to accept other options. The program must still accept the `-data` option with the semantics described above for the release handler to work correctly.

The Erlang emulator is found by examining the registry keys for the emulator version specified in the release data file. The new emulator needs to be properly installed before the upgrade for this to work.

Although the program is located together with files specific to emulator version, it is not expected to be specific to the emulator version. The release handler does *not* change the `-machine` option to `erlsrv` during emulator restart. Place the (possibly customized) `start_erl` program so that it is not overwritten during upgrade.

The `erlsrv` program's default options are not sufficient for release handling. The machine `erlsrv` starts should be specified as the `start_erl` program and the arguments should contain the `++` followed by desired options.

SEE ALSO

`erlsrv(1)`, `release_handler(3)`

werl

Command

On Windows, the preferred way to start the Erlang system for interactive use is:

```
werl <arguments>
```

This will start Erlang in its own window, with fully functioning command-line editing and scrollbars. All flags except `-oldshell` work as they do for the `erl` [page 79] command.

Ctrl-C is reserved for copying text to the clipboard (Ctrl-V to paste). To interrupt the runtime system or the shell process (depending on what has been specified with the `+B` system flag), you should use Ctrl-Break.

In cases where you want to redirect standard input and/or standard output or use Erlang in a pipeline, the `werl` is not suitable, and the `erl` program should be used instead.

The `werl` window is in many ways modelled after the `xterm` window present on other platforms, as the `xterm` model fits well with line oriented command based interaction. This means that selcting text is line oriented rather than rectangle oriented.

To select text in the `werl` window , simply press and hold the left mouse button and drag the mouse over the text you want to select. If the selection crosses line boundaries, the selected text will consist of complete lines where applicable (just like in a word processor). To select more text than fits in the window, start by selecting a small portion in the beginning of the text you want, then use the scrollbar to view the end of the desired selection, point to it and press the *right* mouse-button. The whole area between your first selection and the point where you right-clicked will be included in the selection.

The selected text is copied to the clipboard by either pressing Ctrl-C, using the menu or pressing the copy button in the toolbar.

Pasted text is always inserted at the current prompt position and will be interpreted by Erlang as usual keyboard input.

Previous command lines can be retrieved by pressing the Up arrow or by pressing Ctrl-P. There is also a drop down box in the toolbar containing the command history. Selecting a command in the drop down box will insert it at the prompt, just as if you used the keyboard to retrieve the command.

Closing the `werl` window will stop the Erlang emulator.

erl_set_memory_block

C Module

This documentation is specific to VxWorks.

The `erl_set_memory_block` function/command initiates custom memory allocation for the Erlang emulator. It has to be called before the Erlang emulator is started and makes Erlang use one single large memory block for all memory allocation.

The memory within the block can be utilized by other tasks than Erlang. This is accomplished by calling the functions `sys_alloc`, `sys_realloc` and `sys_free` instead of `malloc`, `realloc` and `free` respectively.

The purpose of this is to avoid problems inherent in the VxWorks systems `malloc` library. The memory allocation within the large memory block avoids fragmentation by using an "address order first fit" algorithm. Another advantage of using a separate memory block is that resource reclamation can be made more easily when Erlang is stopped.

The `erl_set_memory_block` function is callable from any C program as an ordinary 10 argument function as well as from the commandline.

Exports

```
int erl_set_memory_block(size_t size, void *ptr, int warn_mixed_malloc, int
    realloc_always_moves, int use_reclaim, ...)
```

The function is called before Erlang is started to specify a large memory block where Erlang can maintain memory internally.

Parameters:

size_t size The size in bytes of Erlang's internal memory block. Has to be specified. Note that the VxWorks system uses dynamic memory allocation heavily, so leave some memory to the system.

void *ptr A pointer to the actual memory block of size `size`. If this is specified as 0 (NULL), Erlang will allocate the memory when starting and will reclaim the memory block (as a whole) when stopped.

If a memory block is allocated and provided here, the `sys_alloc` etc routines can still be used after the Erlang emulator is stopped. The Erlang emulator can also be restarted while other tasks using the memory block are running without destroying the memory. If Erlang is to be restarted, also set the `use_reclaim` flag.

If 0 is specified here, the Erlang system should not be stopped while some other task uses the memory block (has called `sys_alloc`).

int warn_mixed_malloc If this flag is set to true (anything else than 0), the system will write a warning message on the console if a program is mixing normal `malloc` with `sys_realloc` or `sys_free`.

int realloc_always_moves If this flag is set to true (anything else than 0), all calls to `sys_realloc` result in a moved memory block. This can in certain conditions give less fragmentation. This flag may be removed in future releases.

int use_reclaim If this flag is set to true (anything else than 0), all memory allocated with `sys_alloc` is automatically reclaimed as soon as a task exits. This is very useful to make writing port programs (and other programs as well) easier. Combine this with using the routines `save_open` etc. specified in the `reclaim.h` file delivered in the Erlang distribution.

Return Value:

Returns 0 (OK) on success, otherwise a value $\neq 0$.

```
int erl_memory_show(...)
```

Return Value:

Returns 0 (OK) on success, otherwise a value $\neq 0$.

```
int erl_mem_info_get(MEM_PART_STATS *stats)
```

Parameter:

MEM_PART_STATS *stats A pointer to a `MEM_PART_STATS` structure as defined in `<memLib.h>`. A successful call will fill in all fields of the structure, on error all fields are left untouched.

Return Value:

Returns 0 (OK) on success, otherwise a value $\neq 0$

NOTES

The memory block used by Erlang actually does not need to be inside the area known to ordinary `malloc`. It is possible to set the `USER_RESERVED_MEM` preprocessor symbol when compiling the wind kernel and then use user reserved memory for Erlang. Erlang can therefor utilize memory above the 32 Mb limit of VxWorks on the PowerPC architecture.

Example:

In `config.h` for the wind kernel:

```
#undef LOCAL_MEM_AUTOSIZE
#undef LOCAL_MEM_SIZE
#undef USER_RESERVED_MEM

#define LOCAL_MEM_SIZE      0x05000000
#define USER_RESERVED_MEM  0x03000000
```

In the start-up script/code for the VxWorks node:

```
erl_set_memory_block(sysPhysMemTop()-sysMemTop(), sysMemTop(), 0, 0, 1);
```

Setting the `use_reclaim` flag decreases performance of the system, but makes programming much easier. Other similar facilities are present in the Erlang system even without using a separate memory block. The routines called `save_malloc`, `save_realloc` and `save_free` provide the same facilities by using VxWorks own `malloc`. Similar routines exist for files, see the file `reclaim.h` in the distribution.

erts_alloc

C Module

`erts_alloc` is an Erlang Run-Time System internal memory allocator library. `erts_alloc` provides the Erlang Run-Time System with a number of memory allocators.

Allocators

Currently the following allocators are present:

`temp_alloc` Allocator used for temporary allocations.

`eheap_alloc` Allocator used for Erlang heap data, such as Erlang process heaps.

`binary_alloc` Allocator used for Erlang binary data.

`ets_alloc` Allocator used for ETS data.

`driver_alloc` Allocator used for driver data.

`sl_alloc` Allocator used for memory blocks that are expected to be short-lived.

`ll_alloc` Allocator used for memory blocks that are expected to be long-lived, for example Erlang code.

`fix_alloc` A very fast allocator used for some fix-sized data. `fix_alloc` manages a set of memory pools from which memory blocks are handed out. `fix_alloc` allocates memory pools from `ll_alloc`. Memory pools that have been allocated are never deallocated.

`std_alloc` Allocator used for most memory blocks not allocated via any of the other allocators described above.

`sys_alloc` This is normally the default `malloc` implementation used on the specific OS.

`mseg_alloc` A memory segment allocator. `mseg_alloc` is used by other allocators for allocating memory segments and is currently only available on systems that have the `mmap` system call. Memory segments that are deallocated are kept for a while in a segment cache before they are destroyed. When segments are allocated, cached segments are used if possible instead of creating new segments. This in order to reduce the number of system calls made.

`sys_alloc` and `fix_alloc` are always enabled and cannot be disabled. `mseg_alloc` is always enabled if it is available and an allocator that uses it is enabled. All other allocators can be enabled or disabled [page 107]. By default all allocators are enabled. When an allocator is disabled, `sys_alloc` is used instead of the disabled allocator.

The main idea with the `erts_alloc` library is to separate memory blocks that are used differently into different memory areas, and by this achieving less memory fragmentation. By putting less effort in finding a good fit for memory blocks that are frequently allocated than for those less frequently allocated, a performance gain can be achieved.

The alloc_util framework

Internally a framework called `alloc_util` is used for implementing allocators. `sys_alloc`, `fix_alloc`, and `mseg_alloc` do not use this framework; hence, the following does *not* apply to them.

An allocator manages multiple areas, called carriers, in which memory blocks are placed. A carrier is either placed in a separate memory segment (allocated via `mseg_alloc`) or in the heap segment (allocated via `sys_alloc`). Multiblock carriers are used for storage of several blocks. Singleblock carriers are used for storage of one block. Blocks that are larger than the value of the singleblock carrier threshold (`sbct` [page 108]) parameter are placed in singleblock carriers. Blocks smaller than the value of the `sbct` parameter are placed in multiblock carriers. Normally an allocator creates a “main multiblock carrier”. Main multiblock carriers are never deallocated. The size of the main multiblock carrier is determined by the value of the `mmbcs` [page 108] parameter.

Sizes of multiblock carriers allocated via `mseg_alloc` are decided based on the values of the largest multiblock carrier size (`lmbcs` [page 107]), the smallest multiblock carrier size (`smbcs` [page 108]), and the multiblock carrier growth stages (`mbcgs` [page 108]) parameters. If nc is the current number of multiblock carriers (the main multiblock carrier excluded) managed by an allocator, the size of the next `mseg_alloc` multiblock carrier allocated by this allocator will roughly be $smbcs + nc * (lmbcs - smbcs) / mbcgs$ when $nc \leq mbcgs$, and `lmbcs` when $nc > mbcgs$. If the value of the `sbct` parameter should be larger than the value of the `lmbcs` parameter, the allocator may have to create multiblock carriers that are larger than the value of the `lmbcs` parameter, though. Singleblock carriers allocated via `mseg_alloc` are sized to whole pages.

Sizes of carriers allocated via `sys_alloc` are decided based on the value of the `sys_alloc` carrier size (`ycc` [page 108]) parameter. The size of a carrier is the least number of multiples of the value of the `ycc` parameter that satisfies the request.

Coalescing of free blocks are always performed immediately. Boundary tags (headers and footers) in free blocks are used which makes the time complexity for coalescing constant.

The memory allocation strategy used for multiblock carriers by an allocator is configurable via the `as` [page 107] parameter. Currently the following strategies are available:

Best fit Strategy: Find the smallest block that satisfies the requested block size.

Implementation: A balanced binary search tree is used. The time complexity is proportional to $\log N$, where N is the number of sizes of free blocks.

Address order best fit Strategy: Find the smallest block that satisfies the requested block size. If multiple blocks are found, choose the one with the lowest address.

Implementation: A balanced binary search tree is used. The time complexity is proportional to $\log N$, where N is the number of free blocks.

Good fit Strategy: Try to find the best fit, but settle for the best fit found during a limited search.

Implementation: The implementation uses segregated free lists with a maximum block search depth (in each list) in order to find a good fit fast. When the maximum block search depth is small (by default 3) this implementation has a time complexity that is constant. The maximum block search depth is configurable via the `mbsd` [page 107] parameter.

A fit Strategy: Do not search for a fit, inspect only one free block to see if it satisfies the request. This strategy is only intended to be used for temporary allocations.

Implementation: Inspect the first block in a free-list. If it satisfies the request, it is used; otherwise, a new carrier is created. The implementation has a time complexity that is constant.

As of erts version 5.6.1 the emulator will refuse to use this strategy on other allocators than `temp_alloc`. This since it will only cause problems for other allocators.

System Flags Effecting erts_alloc

Warning:

Only use these flags if you are absolutely sure what you are doing. Unsuitable settings may cause serious performance degradation and even a system crash at any time during operation.

Memory allocator system flags have the following syntax: `+M<S><P> <V>` where `<S>` is a letter identifying a subsystem, `<P>` is a parameter, and `<V>` is the value to use. The flags can be passed to the Erlang emulator (`erl` [page 79]) as command line arguments.

System flags effecting specific allocators have an upper-case letter as `<S>`. The following letters are used for the currently present allocators:

- B: `binary_alloc`
- D: `std_alloc`
- E: `ets_alloc`
- F: `fix_alloc`
- H: `eheap_alloc`
- L: `ll_alloc`
- M: `mseg_alloc`
- R: `driver_alloc`
- S: `sl_alloc`
- T: `temp_alloc`
- Y: `sys_alloc`

The following flags are available for configuration of `mseg_alloc`:

- `+MMamcbf <size>` Absolute max cache bad fit (in kilobytes). A segment in the memory segment cache is not reused if its size exceeds the requested size with more than the value of this parameter. Default value is 4096.
- `+MMrmcbf <ratio>` Relative max cache bad fit (in percent). A segment in the memory segment cache is not reused if its size exceeds the requested size with more than relative max cache bad fit percent of the requested size. Default value is 20.
- `+MMmcs <amount>` Max cached segments. The maximum number of memory segments stored in the memory segment cache. Valid range is 0-30. Default value is 5.

+MMcci <time> Cache check interval (in milliseconds). The memory segment cache is checked for segments to destroy at an interval determined by this parameter. Default value is 1000.

The following flags are available for configuration of `fix_alloc`:

+MFe true Enable `fix_alloc`. Note: `fix_alloc` cannot be disabled.

The following flags are available for configuration of `sys_alloc`:

+MYe true Enable `sys_alloc`. Note: `sys_alloc` cannot be disabled.

+MYm libc malloc library to use. Currently only `libc` is available. `libc` enables the standard `libc` malloc implementation. By default `libc` is used.

+MYtt <size> Trim threshold size (in kilobytes). This is the maximum amount of free memory at the top of the heap (allocated by `sbrk`) that will be kept by `malloc` (not released to the operating system). When the amount of free memory at the top of the heap exceeds the trim threshold, `malloc` will release it (by calling `sbrk`). Trim threshold is given in kilobytes. Default trim threshold is 128. *Note:* This flag will only have any effect when the emulator has been linked with the GNU C library, and uses its `malloc` implementation.

+MYtp <size> Top pad size (in kilobytes). This is the amount of extra memory that will be allocated by `malloc` when `sbrk` is called to get more memory from the operating system. Default top pad size is 0. *Note:* This flag will only have any effect when the emulator has been linked with the GNU C library, and uses its `malloc` implementation.

The following flags are available for configuration of allocators based on `alloc_util`. If `u` is used as subsystem identifier (i.e., <S> = `u`) all allocators based on `alloc_util` will be effected. If `B`, `D`, `E`, `H`, `L`, `R`, `S`, or `T` is used as subsystem identifier, only the specific allocator identified will be effected:

+M<S>as bf|aobf|gf|af Allocation strategy. Valid strategies are `bf` (best fit), `aobf` (address order best fit), `gf` (good fit), and `af` (a fit). See the description of allocation strategies [page 105] in “the `alloc_util` framework” section.

+M<S>asbcst <size> Absolute singleblock carrier shrink threshold (in kilobytes). When a block located in an `mseg_alloc` singleblock carrier is shrunk, the carrier will be left unchanged if the amount of unused memory is less than this threshold; otherwise, the carrier will be shrunk. See also `rsbcst` [page 108].

+M<S>e true|false Enable allocator <S>.

+M<S>lmbcs <size> Largest (`mseg_alloc`) multiblock carrier size (in kilobytes). See the description on how sizes for `mseg_alloc` multiblock carriers are decided [page 105] in “the `alloc_util` framework” section.

+M<S>mbcgs <ratio> (`mseg_alloc`) multiblock carrier growth stages. See the description on how sizes for `mseg_alloc` multiblock carriers are decided [page 105] in “the `alloc_util` framework” section.

+M<S>mbsd <depth> Max block search depth. This flag has effect only if the good fit strategy has been selected for allocator <S>. When the good fit strategy is used, free blocks are placed in segregated free-lists. Each free list contains blocks of sizes in a specific range. The max block search depth sets a limit on the maximum number of blocks to inspect in a free list during a search for suitable block satisfying the request.

- +M<S>mmbcs <size> Main multiblock carrier size. Sets the size of the main multiblock carrier for allocator <S>. The main multiblock carrier is allocated via `sys_alloc` and is never deallocated.
- +M<S>mmmbc <amount> Max `mseg_alloc` multiblock carriers. Maximum number of multiblock carriers allocated via `mseg_alloc` by allocator <S>. When this limit has been reached, new multiblock carriers will be allocated via `sys_alloc`.
- +M<S>mmsbc <amount> Max `mseg_alloc` singleblock carriers. Maximum number of singleblock carriers allocated via `mseg_alloc` by allocator <S>. When this limit has been reached, new singleblock carriers will be allocated via `sys_alloc`.
- +M<S>ramv <bool> Realloc always moves. When enabled, reallocate operations will more or less be translated into an allocate, copy, free sequence. This often reduce memory fragmentation, but costs performance.
- +M<S>rsbcmt <ratio> Relative singleblock carrier move threshold (in percent). When a block located in a singleblock carrier is shrunk to a size smaller than the value of the `sbct` [page 108] parameter, the block will be left unchanged in the singleblock carrier if the ratio of unused memory is less than this threshold; otherwise, it will be moved into a multiblock carrier.
- +M<S>rsbcst <ratio> Relative singleblock carrier shrink threshold (in percent). When a block located in an `mseg_alloc` singleblock carrier is shrunk, the carrier will be left unchanged if the ratio of unused memory is less than this threshold; otherwise, the carrier will be shrunk. See also `asbcst` [page 107].
- +M<S>sbct <size> Singleblock carrier threshold. Blocks larger than this threshold will be placed in singleblock carriers. Blocks smaller than this threshold will be placed in multiblock carriers.
- +M<S>smbcs <size> Smallest (`mseg_alloc`) multiblock carrier size (in kilobytes). See the description on how sizes for `mseg_alloc` multiblock carriers are decided [page 105] in “the `alloc_util` framework” section.
- +M<S>t true|false|<amount> Multiple, thread specific instances of the allocator. This option will only have any effect on the runtime system with SMP support. Default behaviour on the runtime system with SMP support (N equals the number of scheduler threads):
 - `temp_alloc` $N + 1$ instances.
 - `ll_alloc` 1 instance.
 - Other allocators** N instances when N is less than or equal to 16. 16 instances when N is greater than 16.

`temp_alloc` will always use $N + 1$ instances when this option has been enabled regardless of the amount passed. Other allocators will use the same amount of instances as the amount passed as long as it isn't greater than N .

Currently the following flags are available for configuration of `alloc_util`, i.e. all allocators based on `alloc_util` will be effected:

- +Muycs <size> `sys_alloc` carrier size. Carriers allocated via `sys_alloc` will be allocated in sizes which are multiples of the `sys_alloc` carrier size. This is not true for main multiblock carriers and carriers allocated during a memory shortage, though.
- +Mummc <amount> Max `mseg_alloc` carriers. Maximum number of carriers placed in separate memory segments. When this limit has been reached, new carriers will be placed in memory retrieved from `sys_alloc`.

Instrumentation flags:

`+Mim true|false` A map over current allocations is kept by the emulator. The allocation map can be retrieved via the `instrument` module. `+Mim true` implies `+Mis true`. `+Mim true` is the same as `-instr` [page 81].

`+Mis true|false` Status over allocated memory is kept by the emulator. The allocation status can be retrieved via the `instrument` module.

`+Mit X` Reserved for future use. Do *not* use this flag.

Note:

When instrumentation of the emulator is enabled, the emulator uses more memory and runs slower.

Other flags:

`+Mea min|max|r9c|r10b|r11b|config min` Disables all allocators that can be disabled.

`max` Enables all allocators (currently default).

`r9c|r10b|r11b` Configures all allocators as they were configured in respective OTP release. These will eventually be removed.

`config` Disables features that cannot be enabled while creating an allocator configuration with `[erts_alloc_config(3)]`. Note, this option should only be used while running `erts_alloc_config`, *not* when using the created configuration.

Only some default values have been presented here. `[erlang:system_info(allocator)]`, and `[erlang:system_info({allocator, Alloc})]` can be used in order to obtain currently used settings and current status of the allocators.

Note:

Most of these flags are highly implementation dependent, and they may be changed or removed without prior notice.

`erts_alloc` is not obliged to strictly use the settings that have been passed to it (it may even ignore them).

`[erts_alloc_config(3)]` is a tool that can be used to aid creation of an `erts_alloc` configuration that is suitable for a limited number of runtime scenarios.

SEE ALSO

`[erts_alloc_config(3)]`, `erl(1)` [page 79], `[instrument(3)]`, `[erlang(3)]`

driver_entry

Erlang Module

As of erts version 5.5.3 the driver interface has been extended (see extended marker [page 114]). The extended interface introduce version management [page 119], the possibility to pass capability flags (see driver flags [page 115]) to the runtime system at driver initialization, and some new driver API functions.

Note:

Old drivers (compiled with an `erl_driver.h` from an earlier erts version than 5.5.3) have to be recompiled (but does not have to use the extended interface).

The `driver_entry` structure is a C struct that all erlang drivers defines. It contains entry points for the erlang driver that are called by the erlang emulator when erlang code accesses the driver.

The `erl_driver` [page 110] driver API functions needs a port handle that identifies the driver instance (and the port in the emulator). This is only passed to the `start` function, but not to the other functions. The `start` function returns a driver-defined handle that is passed to the other functions. A common practice is to have the `start` function allocating some application-defined structure and stash the port handle in it, to use it later with the driver API functions.

The driver call-back functions are called synchronously from the erlang emulator. If they take too long before completing, they can cause timeouts in the emulator. Use the queue or asynchronous calls if nessecary, since the emulator must be responsive.

The driver structure contains the name of the driver and some 15 function pointers. These pointers are called at different times by the emulator.

The only exported function from the driver is `driver_init`. This function returns the `driver_entry` structure that points to the other functions in the driver. The `driver_init` function is declared with a macro `DRIVER_INIT(drivername)`. (This is because different OS's have different names for it.)

When writing a driver in C++, the driver entry should be of "C" linkage. One way to do this is to put this line somewhere before the driver entry: `extern "C" DRIVER_INIT(drivername);`

When the driver has passed the `driver_entry` over to the emulator, the driver is *not* allowed to modify the `driver_entry`.

Note:

Do *not* declare the `driver_entryconst`. This since the emulator needs to modify the `handle`, and the `handle2` fields. A statically allocated, and `const` declared `driver_entry` may be located in read only memory which will cause the emulator to crash.

Here is the declaration of `driver_entry`:

```
typedef struct erl_drv_entry {
    int (*init)(void);          /* called at system start up for statically
                               linked drivers, and after loading for
                               dynamically loaded drivers */

#ifdef ERL_SYS_DRV
    ErlDrvData (*start)(ErlDrvPort port, char *command);
                               /* called when open_port/2 is invoked.
                               return value -1 means failure. */
#else
    ErlDrvData (*start)(ErlDrvPort port, char *command, SysDriverOpts* opts);
                               /* special options, only for system driver */
#endif
    void (*stop)(ErlDrvData drv_data);
                               /* called when port is closed, and when the
                               emulator is halted. */
    void (*output)(ErlDrvData drv_data, char *buf, int len);
                               /* called when we have output from erlang to
                               the port */
    void (*ready_input)(ErlDrvData drv_data, ErlDrvEvent event);
                               /* called when we have input from one of
                               the driver's handles */
    void (*ready_output)(ErlDrvData drv_data, ErlDrvEvent event);
                               /* called when output is possible to one of
                               the driver's handles */
    char *driver_name;        /* name supplied as command
                               in open_port XXX ? */
    void (*finish)(void);     /* called before unloading the driver -
                               DYNAMIC DRIVERS ONLY */
    void *handle;            /* Reserved -- Used by emulator internally */
    int (*control)(ErlDrvData drv_data, unsigned int command, char *buf,
                  int len, char **rbuf, int rlen);
                               /* "ioctl" for drivers - invoked by
                               port_control/3) */
    void (*timeout)(ErlDrvData drv_data); /* Handling of timeout in driver */
    void (*outputv)(ErlDrvData drv_data, ErlIOVec *ev);
                               /* called when we have output from erlang
                               to the port */
    void (*ready_async)(ErlDrvData drv_data, ErlDrvThreadData thread_data);
    void (*flush)(ErlDrvData drv_data);
                               /* called when the port is about to be
                               closed, and there is data in the
                               driver queue that needs to be flushed
                               before 'stop' can be called */
};
```

```

int (*call)(ErlDrvData drv_data, unsigned int command, char *buf,
            int len, char **rbuf, int rlen, unsigned int *flags);
            /* Works mostly like 'control', a synchronous
            call into the driver. */
void (*event)(ErlDrvData drv_data, ErlDrvEvent event,
              ErlDrvEventData event_data);
            /* Called when an event selected by
            driver_event() has occurred */
int extended_marker; /* ERL_DRV_EXTENDED_MARKER */
int major_version; /* ERL_DRV_EXTENDED_MAJOR_VERSION */
int minor_version; /* ERL_DRV_EXTENDED_MINOR_VERSION */
int driver_flags; /* ERL_DRV_FLAGS */
void *handle2; /* Reserved -- Used by emulator internally */
void (*process_exit)(ErlDrvData drv_data, ErlDrvMonitor *monitor);
            /* Called when a process monitor fires */
} ErlDrvEntry;

```

Exports

```
int init(void)
```

This is called directly after the driver has been loaded by `erl_ddll:load_driver/2`. (Actually when the driver is added to the driver list.) The driver should return 0, or if the driver can't initialize, -1.

```
int start(ErlDrvPort port, char* command)
```

This is called when the driver is instantiated, when `open_port/2` is called. The driver should return a number ≥ 0 or a pointer, or if the driver can't be started, one of three error codes should be returned:

`ERL_DRV_ERROR_GENERAL` - general error, no error code

`ERL_DRV_ERROR_ERRNO` - error with error code in `erl_errno`

`ERL_DRV_ERROR_BADARG` - error, badarg

If an error code is returned, the port isn't started.

```
void stop(ErlDrvData drv_data)
```

This is called when the port is closed, with `port_close/1` or `Port ! {self(), close}`. Note that terminating the port owner process also closes the p port.

```
void output(ErlDrvData drv_data, char *buf, int len)
```

This is called when an erlang process has sent data to the port. The data is pointed to by `buf`, and is `len` bytes. Data is sent to the port with `Port ! {self(), {command, Data}}`, or with `port_command/2`. Depending on how the port was opened, it should be either a list of integers 0...255 or a binary. See `open_port/3` and `port_command/2`.

```
void ready_input(ErlDrvData drv_data, ErlDrvEvent event)
```

```
void ready_output(ErlDrvData drv_data, ErlDrvEvent event)
```

This is called when a driver event (given in the `event` parameter) is signaled. This is used to help asynchronous drivers “wake up” when something happens.

On unix the `event` is a pipe or socket handle (or something that the `select` system call understands).

On Windows the `event` is an Event or Semaphore (or something that the `WaitForMultipleObjects` API function understands). (Some trickery in the emulator allows more than the built-in limit of 64 Events to be used.)

To use this with threads and asynchronous routines, create a pipe on unix and an Event on Windows. When the routine completes, write to the pipe (use `SetEvent` on Windows), this will make the emulator call `ready_input` or `ready_output`.

`char *driver_name`

This is the name of the driver, it must correspond to the atom used in `open_port`, and the name of the driver library file (without the extension).

`void finish(void)`

This function is called by the `erl_ddll` driver when the driver is unloaded. (It is only called in dynamic drivers.)

The driver is only unloaded as a result of calling `unload_driver/1`, or when the emulator halts.

`void *handle`

This field is reserved for the emulators internal use. The emulator will modify this field; therefore, it is important that the `driver_entry` isn't declared `const`.

`int control(ErlDrvData drv_data, unsigned int command, char *buf, int len, char **rbuf, int rlen)`

This is a special routine invoked with the `erlang` function `port_control/3`. It works a little like an “`ioctl`” for erlang drivers. The data given to `port_control/3` arrives in `buf` and `len`. The driver may send data back, using `*rbuf` and `rlen`.

This is the fastest way of calling a driver and get a response. It won't make any context switch in the erlang emulator, and requires no message passing. It is suitable for calling C function to get faster execution, when erlang is too slow.

If the driver wants to return data, it should return it in `rbuf`. Data is returned different depending on the port control flags (those that are set with `set_port_control_flags` [page 131]).

If the flag is set to `PORT_CONTROL_FLAG_BINARY`, then `*rbuf` must point to a binary allocated with `driver_alloc_binary` [page 126] or be `NULL`. This binary will be freed automatically after `control` has returned. The driver can retain the returned binary for *read only* access with `driver_binary_inc_refc` [page 127] to be freed later with `driver_free_binary` [page 126]. It is never allowed to alter the binary after `control` has returned. If `*rbuf` is set to `NULL`, an empty list will be returned.

If the flag is set to 0, `rbuf` points to a `char*` containing data, that is returned as a list of integers. When `control` is called, `*rbuf` points to a buffer of `rlen` bytes, which can be used to return data. A larger buffer can be allocated with `driver_alloc` [page 126]. The buffer will be freed automatically after `control` has returned.

Using binaries is faster if more than a few bytes are returned.

The return value is the number of bytes returned in `*rbuf`.

```
void timeout(ErlDrvData drv_data)
```

This function is called any time after the driver's timer reaches 0. The timer is activated with `driver_set_timer`. There are no priorities or ordering among drivers, so if several drivers time out at the same time, any one of them is called first.

```
void outputv(ErlDrvData drv_data, ErlIOVec *ev)
```

This function is called whenever the port is written to. If it is NULL, the output function is called instead. This function is faster than `output`, because it takes an `ErlIOVec` directly, which requires no copying of the data. The port should be in binary mode, see `open_port/2`.

The `ErlIOVec` contains both a `SysIOVec`, suitable for `writenv`, and one or more binaries. If these binaries should be retained, when the driver returns from `outputv`, they can be queued (using `driver_enq_bin` [page 128] for instance), or if they are kept in a static or global variable, the reference counter can be incremented.

```
void ready_async(ErlDrvData drv_data, ErlDrvThreadData thread_data)
```

This function is called after an asynchronous call has completed. The asynchronous call is started with `driver_async` [page 134]. This function is called from the erlang emulator thread, as opposed to the asynchronous function, which is called in some thread (if multithreading is enabled).

```
int call(ErlDrvData drv_data, unsigned int command, char *buf, int len, char **rbuf,
        int rlen, unsigned int *flags)
```

This function is called from `erlang:port_call/3`. It works a lot like the control call-back, but uses the external term format for input and output.

`command` is an integer, obtained from the call from erlang (the second argument to `erlang:port_call/3`).

`buf` and `len` provide the arguments to the call (the third argument to `erlang:port_call/3`). They can be decoded using `ei` functions.

`rbuf` points to a return buffer, `rlen` bytes long. The return data should be a valid erlang term in the external (binary) format. This is converted to an erlang term and returned by `erlang:port_call/3` to the caller. If more space than `rlen` bytes is needed to return data, `*rbuf` can be set to memory allocated with `driver_alloc`. This memory will be freed automatically after `call` has returned.

The return value is the number of bytes returned in `*rbuf`. If `ERL_DRV_ERROR_GENERAL` is returned (or in fact, anything < 0), `erlang:port_call/3` will throw a `BAD_ARG`.

```
void event(ErlDrvData drv_data, ErlDrvEvent event, ErlDrvEventData event_data)
```

Intentionally left undocumented.

```
int extended_marker
```

This field should either be equal to `ERL_DRV_EXTENDED_MARKER` or 0. An old driver (not aware of the extended driver interface) should set this field to 0. If this field is equal to 0, all the fields following this field also *have* to be 0, or NULL in case it is a pointer field.

`int major_version`

This field should equal `ERL_DRV_EXTENDED_MAJOR_VERSION` if the `extended_marker` field equals `ERL_DRV_EXTENDED_MARKER`.

`int minor_version`

This field should equal `ERL_DRV_EXTENDED_MINOR_VERSION` if the `extended_marker` field equals `ERL_DRV_EXTENDED_MARKER`.

`int driver_flags`

This field is used to pass driver capability information to the runtime system. If the `extended_marker` field equals `ERL_DRV_EXTENDED_MARKER`, it should contain 0 or driver flags (`ERL_DRV_FLAG_*`) ored bitwise. Currently the following driver flags exist:

`ERL_DRV_FLAG_USE_PORT_LOCKING` The runtime system will use port level locking on all ports executing this driver instead of driver level locking when the driver is run in a runtime system with SMP support. For more information see the `erl_driver` [page 116] documentation.

`void *handle2`

This field is reserved for the emulators internal use. The emulator will modify this field; therefore, it is important that the `driver_entry` isn't declared `const`.

`void process_exit(ErlDrvData drv_data, ErlDrvMonitor *monitor)`

This callback is called when a monitored process exits. The `drv_data` is the data associated with the port for which the process is monitored (using `driver_monitor_process` [page 130]) and the `monitor` corresponds to the `ErlDrvMonitor` structure filled in when creating the monitor. The driver interface function `driver_get_monitored_process` [page 130] can be used to retrieve the process id of the exiting process as an `ErlDrvTermData`

SEE ALSO

`erl_driver(3)` [page 116], `[erl_dll(3)]`, `[erlang(3)]`, `kernel(3)`

erl_driver

Erlang Module

As of erts version 5.5.3 the driver interface has been extended (see extended marker [page 114]). The extended interface introduce version management [page 119], the possibility to pass capability flags (see driver flags [page 115]) to the runtime system at driver initialization, and some new driver API functions.

Note:

Old drivers (compiled with an `erl_driver.h` from an earlier erts version than 5.5.3) have to be recompiled (but does not have to use the extended interface).

The driver calls back to the emulator, using the API functions declared in `erl_driver.h`. They are used for outputting data from the driver, using timers, etc.

A driver is a library with a set of function that the emulator calls, in response to Erlang functions and message sending. There may be multiple instances of a driver, each instance is connected to an Erlang port. Every port has a port owner process. Communication with the port is normally done through the port owner process.

Most of the functions takes the `port` handle as an argument. This identifies the driver instance. Note that this port handle must be stored by the driver, it is not given when the driver is called from the emulator (see `driver_entry` [page 110]).

Some of the functions takes a parameter of type `ErlDrvBinary`, a driver binary. It should be both allocated and freed by the caller. Using a binary directly avoid one extra copying of data.

Many of the output functions has a “header buffer”, with `hbuf` and `hlen` parameters. This buffer is sent as a list before the binary (or list, depending on port mode) that is sent. This is convenient when matching on messages received from the port. (Although in the latest versions of Erlang, there is the binary syntax, that enables you to match on the beginning of a binary.)

In the runtime system with SMP support, drivers are locked either on driver level or port level (driver instance level). By default driver level locking will be used, i.e., only one emulator thread will execute code in the driver at a time. If port level locking is used, multiple emulator threads may execute code in the driver at the same time. There will only be one thread at a time calling driver call-backs corresponding to the same port, though. In order to enable port level locking set the `ERL_DRV_FLAG_USE_PORT_LOCKING` driver flag [page 115] in the `driver_entry` [page 110] used by the driver. When port level locking is used it is the responsibility of the driver writer to synchronize all accesses to data shared by the ports (driver instances).

Most drivers written before the runtime system with SMP support existed will be able to run in the runtime system with SMP support without being rewritten if driver level locking is used.

Note:

It is assumed that drivers does not access other drivers. If drivers should access each other they have to provide their own mechanism for thread safe synchronization. Such “inter driver communication” is strongly discouraged.

Previously, in the runtime system without SMP support, specific driver call-backs were always called from the same thread. This is *not* the case in the runtime system with SMP support. Regardless of locking scheme used, calls to driver call-backs may be made from different threads, e.g., two consecutive calls to exactly the same call-back for exactly the same port may be made from two different threads. This will for *most* drivers not be a problem, but it might. Drivers that depend on all call-backs being called in the same thread, *have* to be rewritten before being used in the runtime system with SMP support.

Note:

Regardless of locking scheme used, calls to driver call-backs may be made from different threads.

Most functions in this API are *not* thread-safe, i.e., they may *not* be called from an arbitrary thread. Function that are not documented as thread-safe may only be called from driver call-backs or function calls descending from a driver call-back call. Note that driver call-backs may be called from different threads. This, however, is not a problem for any functions in this API, since the emulator have control over these threads.

Note:

Functions not explicitly documented as thread-safe are *not* thread-safe. Also note that some functions are *only* thread safe when used in a runtime system with SMP support.

Functionality

All functions that a driver needs to do with Erlang are performed through driver API functions. There are functions for the following functionality:

Timer functions Timer functions are used to control the timer that a driver may use. The timer will have the emulator call the `timeout` [page 114] entry function after a specified time. Only one timer is available for each driver instance.

Queue handling Every driver instance has an associated queue. This queue is a `SysIOVec` that works as a buffer. It's mostly used for the driver to buffer data that should be written to a device, it is a byte stream. If the port owner process closes the driver, and the queue is not empty, the driver will not be closed. This enables the driver to flush its buffers before closing.

The queue can be manipulated from arbitrary threads if a port data lock is used. See documentation of the `ErlDrvPDL` [page 122] type for more information.

Output functions With the output functions, the driver sends data back the emulator.

They will be received as messages by the port owner process, see `open_port/2`. The vector function and the function taking a driver binary is faster, because that avoid copying the data buffer. There is also a fast way of sending terms from the driver, without going through the binary term format.

Failure The driver can exit and signal errors up to Erlang. This is only for severe errors, when the driver can't possibly keep open.

Asynchronous calls The latest Erlang versions (R7B and later) has provision for asynchronous function calls, using a thread pool provided by Erlang. There is also a `select` call, that can be used for asynchronous drivers.

Multi-threading A POSIX thread like API for multi-threading is provided. The Erlang driver thread API only provide a subset of the functionality provided by the POSIX thread API. The subset provided is more or less the basic functionality needed for multi-threaded programming:

- Threads [page 122]
- Mutexes [page 122]
- Condition variables [page 123]
- Read/Write locks [page 123]
- Thread specific data [page 123]

The Erlang driver thread API can be used in conjunction with the POSIX thread API on UN-ices and with the Windows native thread API on Windows. The Erlang driver thread API has the advantage of being portable, but there might exist situations where you want to use functionality from the POSIX thread API or the Windows native thread API.

The Erlang driver thread API only return error codes when it is reasonable to recover from an error condition. If it isn't reasonable to recover from an error condition, the whole runtime system is terminated. For example, if a create mutex operation fails, an error code is returned, but if a lock operation on a mutex fails, the whole runtime system is terminated.

Note that there exist no "condition variable wait with timeout" in the Erlang driver thread API. This is due to issues with `pthread_cond_timedwait()`. When the system clock suddenly is changed, it isn't always guaranteed that you will wake up from the call as expected. An Erlang runtime system has to be able to cope with sudden changes of the system clock. Therefore, we have omitted it from the Erlang driver thread API. In the Erlang driver case, timeouts can and should be handled with the timer functionality of the Erlang driver API.

In order for the Erlang driver thread API to function, thread support has to be enabled in the runtime system. An Erlang driver can check if thread support is enabled by use of `driver_system_info()` [page 123]. Note that some functions in the Erlang driver API are thread-safe only when the runtime system has SMP support, also this information can be retrieved via `driver_system_info()` [page 123]. Also note that a lot of functions in the Erlang driver API are *not* thread-safe regardless of whether SMP support is enabled or not. If a function isn't documented as thread-safe it is *not* thread-safe.

NOTE: When executing in an emulator thread, it is *very important* that you unlock *all* locks you have locked before letting the thread out of your control; otherwise, you are *very likely* to deadlock the whole emulator. If you need to use thread specific data in an emulator thread, only have the thread specific data set while the thread is under your control, and clear the thread specific data before you let the thread out of your control.

In the future there will probably be debug functionality integrated with the Erlang driver thread API. All functions that create entities take a `name` argument. Currently the `name` argument is unused, but it will be used when the debug functionality has been implemented. If you name all entities created well, the debug functionality will be able to give you better error reports.

Adding / remove drivers A driver can add and later remove drivers.

Monitoring processes A driver can monitor a process that does not own a port.

Version management Version management is enabled for drivers that have set the `extended_marker` [page 114] field of their `driver_entry` [page 110] to `ERL_DRV_EXTENDED_MARKER`. `erl_driver.h` defines `ERL_DRV_EXTENDED_MARKER`, `ERL_DRV_EXTENDED_MAJOR_VERSION`, and `ERL_DRV_EXTENDED_MINOR_VERSION`. `ERL_DRV_EXTENDED_MAJOR_VERSION` will be incremented when driver incompatible changes are made to the Erlang runtime system. Normally it will suffice to recompile drivers when the `ERL_DRV_EXTENDED_MAJOR_VERSION` has changed, but it could, under rare circumstances, mean that drivers have to be slightly modified. If so, this will of course be documented. `ERL_DRV_EXTENDED_MINOR_VERSION` will be incremented when new features are added. The runtime system use the minor version of the driver to determine what features to use. The runtime system will refuse to load a driver if the major versions differ, or if the major versions are equal and the minor version used by the driver is greater than the one used by the runtime system.

The emulator tries to check that a driver that doesn't use the extended driver interface isn't incompatible when loading it. It can, however, not make sure that it isn't incompatible. Therefore, when loading a driver that doesn't use the extended driver interface, there is a risk that it will be loaded also when the driver is incompatible. When the driver use the extended driver interface, the emulator can verify that it isn't of an incompatible driver version. You are therefore advised to use the extended driver interface.

Exports

`Er1DrvSysInfo`

Types:

- `int driver_major_version`
- `int driver_minor_version`
- `char *erts_version`
- `char *otp_release`
- `int thread_support`
- `int smp_support`
- `int async_threads`
- `int scheduler_threads`

The `Er1DrvSysInfo` structure is used for storage of information about the Erlang runtime system. `driver_system_info()` [page 123] will write the system information when passed a reference to a `Er1DrvSysInfo` structure. A description of the fields in the structure follow:

- driver_major_version** The value of `ERL_DRV_EXTENDED_MAJOR_VERSION` [page 119] when the runtime system was compiled. This value is the same as the value of `ERL_DRV_EXTENDED_MAJOR_VERSION` [page 119] used when compiling the driver; otherwise, the runtime system would have refused to load the driver.
- driver_minor_version** The value of `ERL_DRV_EXTENDED_MINOR_VERSION` [page 119] when the runtime system was compiled. This value might differ from the value of `ERL_DRV_EXTENDED_MINOR_VERSION` [page 119] used when compiling the driver.
- erts_version** A string containing the version number of the runtime system (the same as returned by `[erlang:system_info(version)]`).
- otp_release** A string containing the OTP release number (the same as returned by `[erlang:system_info(otp_release)]`).
- thread_support** A value `!= 0` if the runtime system has thread support; otherwise, 0.
- smp_support** A value `!= 0` if the runtime system has SMP support; otherwise, 0.
- thread_support** A value `!= 0` if the runtime system has thread support; otherwise, 0.
- smp_support** A value `!= 0` if the runtime system has SMP support; otherwise, 0.
- async_threads** The number of async threads in the async thread pool used by `driver_async()` [page 134] (the same as returned by `[erlang:system_info(thread_pool_size)]`).
- scheduler_threads** The number of scheduler threads used by the runtime system (the same as returned by `[erlang:system_info(schedulers)]`).

ErlDrvBinary

Types:

- `int orig_size`
- `char orig_bytes[]`

The `ErlDrvBinary` structure is a binary, as sent between the emulator and the driver. All binaries are reference counted; when `driver_binary_free` is called, the reference count is decremented, when it reaches zero, the binary is deallocated. The `orig_size` is the size of the binary, and `orig_bytes` is the buffer. The `ErlDrvBinary` does not have a fixed size, its size is `orig_size + 2 * sizeof(int)`.

Note:

The `refc` field has been removed. The reference count of an `ErlDrvBinary` is now stored elsewhere. The reference count of an `ErlDrvBinary` can be accessed via `driver_binary_get_refc()` [page 127], `driver_binary_inc_refc()` [page 127], and `driver_binary_dec_refc()` [page 127].

Some driver calls, such as `driver_enq_binary`, increments the driver reference count, and others, such as `driver_deq` decrements it.

Using a driver binary instead of a normal buffer, is often faster, since the emulator doesn't need to copy the data, only the pointer is used.

A driver binary allocated in the driver, with `driver_alloc_binary`, should be freed in the driver (unless otherwise stated), with `driver_free_binary`. (Note that this doesn't

necessarily deallocate it, if the driver is still referred in the emulator, the ref-count will not go to zero.)

Driver binaries are used in the `driver_output2` and `driver_outputv` calls, and in the queue. Also the driver call-back `outputv` [page 114] uses driver binaries.

If the driver of some reason or another, wants to keep a driver binary around, in a static variable for instance, the reference count should be incremented, and the binary can later be freed in the `stop` [page 112] call-back, with `driver_free_binary`.

Note that since a driver binary is shared by the driver and the emulator, a binary received from the emulator or sent to the emulator, must not be changed by the driver.

From erts version 5.5 (OTP release R11B), `orig_bytes` is guaranteed to be properly aligned for storage of an array of doubles (usually 8-byte aligned).

`ErlDrvData`

The `ErlDrvData` is a handle to driver-specific data, passed to the driver call-backs. It is a pointer, and is most often casted to a specific pointer in the driver.

`SysIOVec`

This is a system I/O vector, as used by `writv` on unix and `WSASend` on Win32. It is used in `ErllIOVec`.

`ErllIOVec`

Types:

- `int vsize`
- `int size`
- `SysIOVec* iov`
- `ErlDrvBinary** binv`

The I/O vector used by the emulator and drivers, is a list of binaries, with a `SysIOVec` pointing to the buffers of the binaries. It is used in `driver_outputv` and the `outputv` [page 114] driver call-back. Also, the driver queue is an `ErllIOVec`.

`ErlDrvMonitor`

When a driver creates a monitor for a process, a `ErlDrvMonitor` is filled in. This is an opaque data-type which can be assigned to but not compared without using the supplied compare function (i.e. it behaves like a struct).

The driver writer should provide the memory for storing the monitor when calling `driver_monitor_process` [page 130]. The address of the data is not stored outside of the driver, so the `ErlDrvMonitor` can be used as any other datum, it can be copied, moved in memory, forgotten etc.

`ErlDrvNowData`

The `ErlDrvNowData` structure holds a timestamp consisting of three values measured from some arbitrary point in the past. The three structure members are:

megasecs The number of whole megaseconds elapsed since the arbitrary point in time

secs The number of whole seconds elapsed since the arbitrary point in time

microsecs The number of whole microseconds elapsed since the arbitrary point in time

ErlDrvPDL

If certain port specific data have to be accessed from other threads than those calling the driver call-backs, a port data lock can be used in order to synchronize the operations on the data. Currently, the only port specific data that the the emulator associates with the port data lock is the driver queue.

Normally a driver instance does not have a port data lock. If the driver instance want to use a port data lock, it has to create the port data lock by calling `driver_pdl_create()` [page 129]. *NOTE:* Once the port data lock has been created, every access to data associated with the port data lock have to be done while having the port data lock locked. The port data lock is locked, and unlocked, respectively, by use of `driver_pdl_lock()` [page 129], and `driver_pdl_unlock()` [page 129].

A port data lock is reference counted, and when the reference count reach zero, it will be destroyed. The emulator will at least increment the reference count once when the lock is created and decrement it once when the port associated with the lock terminates. The emulator will also increment the reference count when an async job is enqueued and decrement it after an async job has been invoked, or canceled. Besides this, it is the responsibility of the driver to ensure that the reference count does not reach zero before the last use of the lock by the driver has been made. The reference count can be read, incremented, and decremented, respectively, by use of `driver_pdl_get_refc()` [page 129], `driver_pdl_inc_refc()` [page 129], and `driver_pdl_dec_refc()` [page 129].

ErlDrvTid

Thread identifier.

See also: `erl_drv_thread_create()` [page 136], `erl_drv_thread_exit()` [page 137], `erl_drv_thread_join()` [page 137], `erl_drv_thread_self()` [page 138], and `erl_drv_equal_tids()` [page 138].

ErlDrvThreadOpts

Types:

- `int suggested_stack_size`

Thread options structure passed to `erl_drv_thread_create()` [page 136]. Currently the following fields exist:

suggested_stack_size A suggestion, in kilo-words, on how large stack to use. A value less than zero means default size.

See also: `erl_drv_thread_opts_create()` [page 137], `erl_drv_thread_opts_destroy()` [page 137], and `erl_drv_thread_create()` [page 136].

ErlDrvMutex

Mutual exclusion lock. Used for synchronizing access to shared data. Only one thread at a time can lock a mutex.

See also: `erl_drv_mutex_create()` [page 138], `erl_drv_mutex_destroy()` [page 138], `erl_drv_mutex_lock()` [page 139], `erl_drv_mutex_trylock()` [page 139], and `erl_drv_mutex_unlock()` [page 139].

ErlDrvCond

Condition variable. Used when threads need to wait for a specific condition to appear before continuing execution. Condition variables need to be used with associated mutexes.

See also: `erl_drv_cond_create()` [page 139], `erl_drv_cond_destroy()` [page 140], `erl_drv_cond_signal()` [page 140], `erl_drv_cond_broadcast()` [page 140], and `erl_drv_cond_wait()` [page 140].

ErlDrvRWLock

Read/write lock. Used to allow multiple threads to read shared data while only allowing one thread to write the same data. Multiple threads can read lock an rwlock at the same time, while only one thread can read/write lock an rwlock at a time.

See also: `erl_drv_rwlock_create()` [page 141], `erl_drv_rwlock_destroy()` [page 141], `erl_drv_rwlock_rlock()` [page 141], `erl_drv_rwlock_trylock()` [page 142], `erl_drv_rwlock_runlock()` [page 142], `erl_drv_rwlock_rwlock()` [page 142], `erl_drv_rwlock_tryrwlock()` [page 143], and `erl_drv_rwlock_rwunlock()` [page 143].

ErlDrvTSDKey

Key which thread specific data can be associated with.

See also: `erl_drv_tsd_key_create()` [page 143], `erl_drv_tsd_key_destroy()` [page 144], `erl_drv_tsd_set()` [page 144], and `erl_drv_tsd_get()` [page 144].

```
void driver_system_info(ErlDrvSysInfo *sys_info_ptr, size_t size)
```

This function will write information about the Erlang runtime system into the `ErlDrvSysInfo` [page 119] structure referred to by the first argument. The second argument should be the size of the `ErlDrvSysInfo` [page 119] structure, i.e., `sizeof(ErlDrvSysInfo)`.

See the documentation of the `ErlDrvSysInfo` [page 119] structure for information about specific fields.

```
int driver_output(ErlDrvPort port, char *buf, int len)
```

The `driver_output` function is used to send data from the driver up to the emulator. The data will be received as terms or binary data, depending on how the driver port was opened.

The data is queued in the port owner process' message queue. Note that this does not yield to the emulator. (Since the driver and the emulator runs in the same thread.)

The parameter `buf` points to the data to send, and `len` is the number of bytes.

The return value for all output functions is 0. (Unless the driver is used for distribution, in which case it can fail and return -1. For normal use, the output function always returns 0.)

```
int driver_output2(ErlDrvPort port, char *hbuf, int hlen, char *buf, int len)
```

The `driver_output2` function first sends `hbuf` (length in `hlen`) data as a list, regardless of port settings. Then `buf` is sent as a binary or list. E.g. if `hlen` is 3 then the port owner process will receive `[H1, H2, H3 | T]`.

The point of sending data as a list header, is to facilitate matching on the data received.

The return value is 0 for normal use.

```
int driver_output_binary(ErlDrvPort port, char *hbuf, int hlen, ErlDrvBinary* bin, int
    offset, int len)
```

This function sends data to port owner process from a driver binary, it has a header buffer (`hbuf` and `hlen`) just like `driver_output2`. The `hbuf` parameter can be `NULL`. The parameter `offset` is an offset into the binary and `len` is the number of bytes to send.

Driver binaries are created with `driver_alloc_binary`.

The data in the header is sent as a list and the binary as an Erlang binary in the tail of the list.

E.g. if `hlen` is 2, then the port owner process will receive `[H1, H2 | <<T>>]`.

The return value is 0 for normal use.

Note that, using the binary syntax in Erlang, the driver application can match the header directly from the binary, so the header can be put in the binary, and `hlen` can be set to 0.

```
int driver_outputv(ErlDrvPort port, char* hbuf, int hlen, ErlIOVec *ev, int skip)
```

This function sends data from an IO vector, `ev`, to the port owner process. It has a header buffer (`hbuf` and `hlen`), just like `driver_output2`.

The `skip` parameter is a number of bytes to skip of the `ev` vector from the head.

You get vectors of `ErlIOVec` type from the driver queue (see below), and the `outputv` [page 114] driver entry function. You can also make them yourself, if you want to send several `ErlDrvBinary` buffers at once. Often it is faster to use `driver_output` or `driver_output_binary`.

E.g. if `hlen` is 2 and `ev` points to an array of three binaries, the port owner process will receive `[H1, H2, <<B1>>, <<B2>> | <<B3>>]`.

The return value is 0 for normal use.

The comment for `driver_output_binary` applies for `driver_outputv` too.

```
int driver_vec_to_buf(ErlIOVec *ev, char *buf, int len)
```

This function collects several segments of data, referenced by `ev`, by copying them in order to the buffer `buf`, of the size `len`.

If the data is to be sent from the driver to the port owner process, it is faster to use `driver_outputv`.

The return value is the space left in the buffer, i.e. if the `ev` contains less than `len` bytes it's the difference, and if `ev` contains `len` bytes or more, it's 0. This is faster if there is more than one header byte, since the binary syntax can construct integers directly from the binary.

```
int driver_set_timer(ErlDrvPort port, unsigned long time)
```

This function sets a timer on the driver, which will count down and call the driver when it is timed out. The `time` parameter is the time in milliseconds before the timer expires. When the timer reaches 0 and expires, the driver entry function `timeout` [page 110] is called.

Note that there is only one timer on each driver instance; setting a new timer will replace an older one.

Return value is 0 (-1 only when the `timeout` driver function is NULL).

```
int driver_cancel_timer(ErlDrvPort port)
```

This function cancels a timer set with `driver_set_timer`.

The return value is 0.

```
int driver_read_timer(ErlDrvPort port, unsigned long *time_left)
```

This function reads the current time of a timer, and places the result in `time_left`. This is the time in milliseconds, before the timeout will occur.

The return value is 0.

```
int driver_get_now(ErlDrvNowData *now)
```

This function reads a timestamp into the memory pointed to by the parameter `now`. See the description of `ErlDrvNowData` [page 121] for specification of its fields.

The return value is 0 unless the `now` pointer is not valid, in which case it is < 0.

```
int driver_select(ErlDrvPort port, ErlDrvEvent event, int mode, int on)
```

The `driver_select` is used by the driver to provide the emulator with an event to check for. This enables the emulator to call the driver when something has happened asynchronously.

The event parameter is used in the emulator cycle in a `select` call. If the event is set then the driver is called. The `mode` parameter can be either `ON_READ` or `ON_WRITE`, and specifies whether `ready_output` [page 113] or `ready_input` [page 113] will be called when the event is fired. Note that this is just a convention, they don't have to read or write anything.

The `on` parameter should be 1 for adding the event and 0 for removing it.

On Unix systems, the function `select` is used. The event must be a socket or pipe (or other object that `select` can use).

On windows, the Win32 API function `WaitForMultipleObjects` is used. This places other restriction on the event. Refer to the Win32 SDK documentation.

The return value is 0 (Failure, -1, only if the `ready_input/ready_output` is NULL).

```
void *driver_alloc(size_t size)
```

This function allocates a memory block of the size specified in `size`, and returns it. This only fails on out of memory, in that case `NULL` is returned. (This is most often a wrapper for `malloc`).

Memory allocated must be explicitly freed with a corresponding call to `driver_free` (unless otherwise stated).

This function is thread-safe.

```
void *driver_realloc(void *ptr, size_t size)
```

This function resizes a memory block, either in place, or by allocating a new block, copying the data and freeing the old block. A pointer is returned to the reallocated memory. On failure (out of memory), `NULL` is returned. (This is most often a wrapper for `realloc`.)

This function is thread-safe.

```
void driver_free(void *ptr)
```

This function frees the memory pointed to by `ptr`. The memory should have been allocated with `driver_alloc`. All allocated memory should be deallocated, just once. There is no garbage collection in drivers.

This function is thread-safe.

```
ErlDrvBinary* driver_alloc_binary(int size)
```

This function allocates a driver binary with a memory block of at least `size` bytes, and returns a pointer to it, or `NULL` on failure (out of memory). When a driver binary has been sent to the emulator, it must not be altered. Every allocated binary should be freed by a corresponding call to `driver_free_binary` (unless otherwise stated).

Note that a driver binary has an internal reference counter, this means that calling `driver_free_binary` it may not actually dispose of it. If it's sent to the emulator, it may be referenced there.

The driver binary has a field, `orig_bytes`, which marks the start of the data in the binary.

This function is thread-safe.

```
ErlDrvBinary* driver_realloc_binary(ErlDrvBinary *bin, int size)
```

This function resizes a driver binary, while keeping the data. The resized driver binary is returned. On failure (out of memory), `NULL` is returned.

This function is only thread-safe when the emulator with SMP support is used.

```
void driver_free_binary(ErlDrvBinary *bin)
```

This function frees a driver binary `bin`, allocated previously with `driver_alloc_binary`. Since binaries in Erlang are reference counted, the binary may still be around.

This function is only thread-safe when the emulator with SMP support is used.

```
long driver_binary_get_refc(ErlDrvBinary *bin)
```

Returns current reference count on `bin`.

This function is only thread-safe when the emulator with SMP support is used.

```
long driver_binary_inc_refc(ErlDrvBinary *bin)
```

Increments the reference count on `bin` and returns the reference count reached after the increment.

This function is only thread-safe when the emulator with SMP support is used.

```
long driver_binary_dec_refc(ErlDrvBinary *bin)
```

Decrements the reference count on `bin` and returns the reference count reached after the decrement.

This function is only thread-safe when the emulator with SMP support is used.

Note:

You should normally decrement the reference count of a driver binary by calling `driver_free_binary()` [page 126]. `driver_binary_dec_refc()` does *not* free the binary if the reference count reaches zero. *Only* use `driver_binary_dec_refc()` when you are sure *not* to reach a reference count of zero.

```
int driver_enq(ErlDrvPort port, char* buf, int len)
```

This function enqueues data in the driver queue. The data in `buf` is copied (`len` bytes) and placed at the end of the driver queue. The driver queue is normally used in a FIFO way.

The driver queue is available to queue output from the emulator to the driver (data from the driver to the emulator is queued by the emulator in normal erlang message queues). This can be useful if the driver has to wait for slow devices etc, and wants to yield back to the emulator. The driver queue is implemented as an `ErlIOVec`.

When the queue contains data, the driver won't close, until the queue is empty.

The return value is 0.

This function can be called from an arbitrary thread if a port data lock [page 122] associated with the `port` is locked by the calling thread during the call.

```
int driver_pushq(ErlDrvPort port, char* buf, int len)
```

This function puts data at the head of the driver queue. The data in `buf` is copied (`len` bytes) and placed at the beginning of the queue.

The return value is 0.

This function can be called from an arbitrary thread if a port data lock [page 122] associated with the `port` is locked by the calling thread during the call.

```
int driver_deq(ErlDrvPort port, int size)
```

This function dequeues data by moving the head pointer forward in the driver queue by `size` bytes. The data in the queue will be deallocated.

The return value is 0.

This function can be called from an arbitrary thread if a port data lock [page 122] associated with the port is locked by the calling thread during the call.

```
int driver_sizeq(ErlDrvPort port)
```

This function returns the number of bytes currently in the driver queue.

This function can be called from an arbitrary thread if a port data lock [page 122] associated with the port is locked by the calling thread during the call.

```
int driver_enq_bin(ErlDrvPort port, ErlDrvBinary *bin, int offset, int len)
```

This function enqueues a driver binary in the driver queue. The data in `bin` at `offset` with length `len` is placed at the end of the queue. This function is most often faster than `driver_enq`, because the data doesn't have to be copied.

This function can be called from an arbitrary thread if a port data lock [page 122] associated with the port is locked by the calling thread during the call.

The return value is 0.

```
int driver_pushq_bin(ErlDrvPort port, ErlDrvBinary *bin, int offset, int len)
```

This function puts data in the binary `bin`, at `offset` with length `len` at the head of the driver queue. It is most often faster than `driver_pushq`, because the data doesn't have to be copied.

This function can be called from an arbitrary thread if a port data lock [page 122] associated with the port is locked by the calling thread during the call.

The return value is 0.

```
SysIOVec* driver_peekq(ErlDrvPort port, int *vlen)
```

This function retrieves the driver queue as a pointer to an array of `SysIOVecs`. It also returns the number of elements in `vlen`. This is the only way to get data out of the queue.

Nothing is remove from the queue by this function, that must be done with `driver_deq`.

The returned array is suitable to use with the Unix system call `writev`.

This function can be called from an arbitrary thread if a port data lock [page 122] associated with the port is locked by the calling thread during the call.

```
int driver_enqv(ErlDrvPort port, ErlIOVec *ev, int skip)
```

This function enqueues the data in `ev`, skipping the first `skip` bytes of it, at the end of the driver queue. It is faster than `driver_enq`, because the data doesn't have to be copied.

The return value is 0.

This function can be called from an arbitrary thread if a port data lock [page 122] associated with the port is locked by the calling thread during the call.

```
int driver_pushqv(ErlDrvPort port, ErlIOVec *ev, int skip)
```

This function puts the data in `ev`, skipping the first `skip` bytes of it, at the head of the driver queue. It is faster than `driver_pushq`, because the data doesn't have to be copied. The return value is 0.

This function can be called from an arbitrary thread if a port data lock [page 122] associated with the port is locked by the calling thread during the call.

```
ErlDrvPDL driver_pdl_create(ErlDrvPort port)
```

This function creates a port data lock associated with the port. *NOTE:* Once a port data lock has been created, it has to be locked during all operations on the driver queue of the port.

On success a newly created port data lock is returned. On failure NULL is returned. `driver_pdl_create()` will fail if port is invalid or if a port data lock already has been associated with the port.

```
void driver_pdl_lock(ErlDrvPDL pdl)
```

This function locks the port data lock passed as argument (`pdl`).

This function is thread-safe.

```
void driver_pdl_unlock(ErlDrvPDL pdl)
```

This function unlocks the port data lock passed as argument (`pdl`).

This function is thread-safe.

```
long driver_pdl_get_refc(ErlDrvPDL pdl)
```

This function returns the current reference count of the port data lock passed as argument (`pdl`).

This function is thread-safe.

```
long driver_pdl_inc_refc(ErlDrvPDL pdl)
```

This function increments the reference count of the port data lock passed as argument (`pdl`).

The current reference count after the increment has been performed is returned.

This function is thread-safe.

```
long driver_pdl_dec_refc(ErlDrvPDL pdl)
```

This function decrements the reference count of the port data lock passed as argument (`pdl`).

The current reference count after the decrement has been performed is returned.

This function is thread-safe.

```
int driver_monitor_process(ErlDrvPort port, ErlDrvTermData process, ErlDrvMonitor *monitor)
```

Start monitoring a process from a driver. When a process is monitored, a process exit will result in a call to the provided `process_exit` [page 115] call-back in the `ErlDrvEntry` [page 110] structure. The `ErlDrvMonitor` structure is filled in, for later removal or compare.

The `process` parameter should be the return value of an earlier call to `driver_caller` [page 131] or `driver_connected` [page 131] call.

The function returns 0 on success, < 0 if no call-back is provided and > 0 if the process is no longer alive.

```
int driver_demonitor_process(ErlDrvPort port, const ErlDrvMonitor *monitor)
```

This function cancels an monitor created earlier.

The function returns 0 if a monitor was removed and > 0 if the monitor did no longer exist.

```
ErlDrvTermData driver_get_monitored_process(ErlDrvPort port, const ErlDrvMonitor *monitor)
```

The function returns the process id associated with a living monitor. It can be used in the `process_exit` call-back to get the process identification for the exiting process.

The function returns `driver_term_nil` if the monitor no longer exists.

```
int driver_compare_monitors(const ErlDrvMonitor *monitor1, const ErlDrvMonitor *monitor2)
```

This function is used to compare two `ErlDrvMonitors`. It can also be used to imply some artificial order on monitors, for whatever reason.

The function returns 0 if `monitor1` and `monitor2` are equal, < 0 if `monitor1` is less than `monitor2` and > 0 if `monitor1` is greater than `monitor2`.

```
void add_driver_entry(ErlDrvEntry *de)
```

This function adds a driver entry to the list of drivers known by Erlang. The `init` [page 112] function of the `de` parameter is called.

Note:

To use this function for adding drivers residing in dynamically loaded code is dangerous. If the driver code for the added driver resides in the same dynamically loaded module (i.e. `.so` file) as a normal dynamically loaded driver (loaded with the `erl_ddll` interface), the caller should call `driver_lock_driver` [page 135] before adding driver entries.

Use of this function is generally deprecated.

```
int remove_driver_entry(ErlDrvEntry *de)
```

This function removes a driver entry `de` previously added with `add_driver_entry`.

Driver entries added by the `erl_ddll` erlang interface can not be removed by using this interface.


```
char* erl_errno_id(int error)
```

This function returns the atom name of the erlang error, given the error number in error. Error atoms are: `EINVAL`, `ENOENT`, etc. It can be used to make error terms from the driver.

```
void set_busy_port(ErlDrvPort port, int on)
```

This function set and resets the busy status of the port. If `on` is 1, the port is set to busy, if it's 0 the port is set to not busy.

When the port is busy, sending to it with `Port ! Data` or `port_command/2`, will block the port owner process, until the port is signaled as not busy.

```
void set_port_control_flags(ErlDrvPort port, int flags)
```

This function sets flags for how the control [page 113] driver entry function will return data to the port owner process. (The `control` function is called from `port_control/3` in erlang.)

Currently there are only two meaningful values for `flags`: 0 means that data is returned in a list, and `PORT_CONTROL_FLAG_BINARY` means data is returned as a binary from `control`.

```
int driver_failure_eof(ErlDrvPort port)
```

This function signals to erlang that the driver has encountered an EOF and should be closed, unless the port was opened with the `eof` option, in that case `eof` is sent to the port. Otherwise, the port is close and an 'EXIT' message is sent to the port owner process.

The return value is 0.

```
int driver_failure_atom(ErlDrvPort port, char *string)
```

```
int driver_failure_posix(ErlDrvPort port, int error)
```

```
int driver_failure(ErlDrvPort port, int error)
```

These functions signal to Erlang that the driver has encountered an error and should be closed. The port is closed and the tuple `{'EXIT', error, Err}`, is sent to the port owner process, where `error` is an error atom (`driver_failure_atom` and `driver_failure_posix`), or an integer (`driver_failure`).

The driver should fail only when in severe error situations, when the driver cannot possibly keep open, for instance buffer allocation gets out of memory. Normal errors is more appropriate to handle with sending error codes with `driver_output`.

The return value is 0.

```
ErlDrvTermData driver_connected(ErlDrvPort port)
```

This function returns the port owner process.

```
ErlDrvTermData driver_caller(ErlDrvPort port)
```

This function returns the process id of the process that made the current call to the driver. The process id can be used with `driver_send_term` to send back data to the caller. `driver_caller()` only return valid data when currently executing in one of the following driver callbacks:

start [page 112] Called from `open_port/2`.

output [page 112] Called from `erlang:send/2`, and `erlang:port_command/2`

outputv [page 114] Called from `erlang:send/2`, and `erlang:port_command/2`

control [page 113] Called from `erlang:port_control/3`

call [page 114] Called from `erlang:port_call/3`

```
int driver_output_term(ErlDrvPort port, ErlDrvTermData* term, int n)
```

This functions sends data in the special driver term format. This is a fast way to deliver term data from a driver. It also needs no binary conversion, so the port owner process receives data as normal Erlang terms.

The `term` parameter points to an array of `ErlDrvTermData`, with `n` elements. This array contains terms described in the driver term format. Every term consists of one to four elements in the array. The term first has a term type, and then arguments.

Tuple and lists (with the exception of strings, see below), are built in reverse polish notation, so that to build a tuple, the elements are given first, and then the tuple term, with a count. Likewise for lists.

A tuple must be specified with the number of elements. (The elements precedes the `ERL_DRV_TUPLE` term.)

A list must be specified with the number of elements, including the tail, which is the last term preceding `ERL_DRV_LIST`.

The special term `ERL_DRV_STRING_CONS` is used to “splice” in a string in a list, a string given this way is not a list per se, but the elements are elements of the surrounding list.

Term type	Argument(s)
ERL_DRV_NIL	
ERL_DRV_ATOM	<code>ErlDrvTermData atom (from <code>driver_mk_atom(char *string)</code>)</code>
ERL_DRV_INT	<code>ErlDrvSInt integer</code>
ERL_DRV_UINT	<code>ErlDrvUInt integer</code>
ERL_DRV_PORT	<code>ErlDrvTermData port (from <code>driver_mk_port(ErlDrvPort port)</code>)</code>
ERL_DRV_BINARY	<code>ErlDrvBinary *bin, ErlDrvUInt len, ErlDrvUInt offset</code>
ERL_DRV_BUF2BINARY	<code>char *buf, ErlDrvUInt len</code>
ERL_DRV_STRING	<code>char *str, int len</code>
ERL_DRV_TUPLE	<code>int sz</code>
ERL_DRV_LIST	<code>int sz</code>
ERL_DRV_PID	<code>ErlDrvTermData pid (from <code>driver_connected(ErlDrvPort port)</code> or <code>driver_c</code>)</code>
ERL_DRV_STRING_CONS	<code>char *str, int len</code>
ERL_DRV_FLOAT	<code>double *dbl</code>
ERL_DRV_EXT2TERM	<code>char *buf, ErlDrvUInt len</code>

The unsigned integer data type `ErlDrvUInt` and the signed integer data type `ErlDrvSInt` are 64 bits wide on a 64 bit runtime system and 32 bits wide on a 32 bit runtime system. They were introduced in erts version 5.6, and replaced some of the `int` arguments in the list above.

To build the tuple `{tcp, Port, [100 | Binary]}`, the following call could be made.

```

ErlDrvBinary* bin = ...
ErlDrvPort port = ...
ErlDrvTermData spec[] = {
    ERL_DRV_ATOM, driver_mk_atom("tcp"),
    ERL_DRV_PORT, driver_mk_port(port),
    ERL_DRV_INT, 100,
    ERL_DRV_BINARY, bin, 50, 0,
    ERL_DRV_LIST, 2,
    ERL_DRV_TUPLE, 3,
};
driver_output_term(port, spec, sizeof(spec) / sizeof(spec[0]));

```

Where `bin` is a driver binary of length at least 50 and `port` is a port handle. Note that the `ERL_DRV_LIST` comes after the elements of the list, likewise the `ERL_DRV_TUPLE`.

The term `ERL_DRV_STRING_CONS` is a way to construct strings. It works differently from how `ERL_DRV_STRING` works. `ERL_DRV_STRING_CONS` builds a string list in reverse order, (as opposed to how `ERL_DRV_LIST` works), concatenating the strings added to a list. The tail must be given before `ERL_DRV_STRING_CONS`.

The `ERL_DRV_STRING` constructs a string, and ends it. (So it's the same as `ERL_DRV_NIL` followed by `ERL_DRV_STRING_CONS`.)

```

/* to send [x, "abc", y] to the port: */
ErlDrvTermData spec[] = {
    ERL_DRV_ATOM, driver_mk_atom("x"),
    ERL_DRV_STRING, (ErlDrvTermData)"abc", 3,
    ERL_DRV_ATOM, driver_mk_atom("y"),
    ERL_DRV_NIL,
    ERL_DRV_LIST, 4
};
driver_output_term(port, spec, sizeof(spec) / sizeof(spec[0]));

/* to send "abc123" to the port: */
ErlDrvTermData spec[] = {
    ERL_DRV_NIL, /* with STRING_CONS, the tail comes first */
    ERL_DRV_STRING_CONS, (ErlDrvTermData)"123", 3,
    ERL_DRV_STRING_CONS, (ErlDrvTermData)"abc", 3,
};
driver_output_term(port, spec, sizeof(spec) / sizeof(spec[0]));

```

The `ERL_DRV_EXT2TERM` term type is used for passing a term encoded with the external format [page 51], i.e., a term that has been encoded by `[erlang:term_to_binary]`, `[erl_interface]`, etc. For example, if `binp` is a pointer to an `ErlDrvBinary` that contains the term `{17, 4711}` encoded with the external format [page 51] and you want to wrap it in a two tuple with the tag `my_tag`, i.e., `{my_tag, {17, 4711}}`, you can do as follows:

```

ErlDrvTermData spec[] = {
    ERL_DRV_ATOM, driver_mk_atom("my_tag"),
    ERL_DRV_EXT2TERM, (ErlDrvTermData) binp->orig_bytes, binp->orig_size,
    ERL_DRV_TUPLE, 2,
};
driver_output_term(port, spec, sizeof(spec) / sizeof(spec[0]));

```

If you want to pass a binary and doesn't already have the content of the binary in an `ErlDrvBinary`, you can benefit from using `ERL_DRV_BUF2BINARY` instead of creating an `ErlDrvBinary` via `driver_alloc_binary()` and then pass the binary via `ERL_DRV_BINARY`. The runtime system will often allocate binaries smarter if `ERL_DRV_BUF2BINARY` is used. However, if the content of the binary to pass already resides in an `ErlDrvBinary`, it is normally better to pass the binary using `ERL_DRV_BINARY` and the `ErlDrvBinary` in question.

The `ERL_DRV_UINT`, `ERL_DRV_BUF2BINARY`, and `ERL_DRV_EXT2TERM` term types were introduced in the 5.6 version of erts.

Note that this function is *not* thread-safe, not even when the emulator with SMP support is used.

```
ErlDrvTermData driver_mk_atom(char* string)
```

This function returns an atom given a name `string`. The atom is created and won't change, so the return value may be saved and reused, which is faster than looking up the atom several times.

```
ErlDrvTermData driver_mk_port(ErlDrvPort port)
```

This function converts a port handle to the erlang term format, usable in the `driver_output_send` function.

```
int driver_send_term(ErlDrvPort port, ErlDrvTermData receiver, ErlDrvTermData* term,
int n)
```

This function is the only way for a driver to send data to *other* processes than the port owner process. The `receiver` parameter specifies the process to receive the data.

The parameters `term` and `n` does the same thing as in `driver_output_term` [page 132].

This function is only thread-safe when the emulator with SMP support is used.

```
long driver_async (ErlDrvPort port, unsigned int* key, void (*async_invoke)(void*),
void* async_data, void (*async_free)(void*))
```

This function performs an asynchronous call. The function `async_invoke` is invoked in a thread separate from the emulator thread. This enables the driver to perform time-consuming, blocking operations without blocking the emulator.

Erlang is by default started without an async thread pool. The number of async threads that the runtime system should use is specified by the `+A` [page 83] command line argument of `erl(1)` [page 79]. If no async thread pool is available, the call is made synchronously in the thread calling `driver_async()`. The current number of async threads in the async thread pool can be retrieved via `driver_system_info()` [page 123].

If there is a thread pool available, a thread will be used. If the `key` argument is null, the threads from the pool are used in a round-robin way, each call to `driver_async` uses the next thread in the pool. With the `key` argument set, this behaviour is changed. The two same values of `*key` always get the same thread.

To make sure that a driver instance always uses the same thread, the following call can be used:

```
unsigned int myKey = (unsigned int) myPort;

r = driver_async(myPort, &myKey, myData, myFunc);
```

It is enough to initialize `myKey` once for each driver instance.

If a thread is already working, the calls will be queued up and executed in order. Using the same thread for each driver instance ensures that the calls will be made in sequence.

The `async_data` is the argument to the functions `async_invoke` and `async_free`. It's typically a pointer to a structure that contains a pipe or event that can be used to signal that the async operation completed. The data should be freed in `async_free`, because it's called if `driver_async_cancel` is called.

When the async operation is done, `ready_async` [page 114] driver entry function is called. If `async_ready` is null in the driver entry, the `async_free` function is called instead.

The return value is a handle to the asynchronous task, which can be used as argument to `driver_async_cancel`.

Note:

As of erts version 5.5.4.3 the default stack size for threads in the async-thread pool is 16 kilowords, i.e., 64 kilobyte on 32-bit architectures. This small default size has been chosen since the amount of async-threads might be quite large. The default stack size is enough for drivers delivered with Erlang/OTP, but might not be sufficiently large for other dynamically linked in drivers that use the `driver_async()` functionality. A suggested stack size for threads in the async-thread pool can be configured via the `+a` [page 83] command line argument of `erl(1)` [page 79].

```
int driver_async_cancel(long id)
```

This function cancels an asynchronous operation, by removing it from the queue. Only functions in the queue can be cancelled; if a function is executing, it's too late to cancel it. The `async_free` function is also called.

The return value is 1 if the operation was removed from the queue, otherwise 0.

```
int driver_lock_driver(ErlDrvPort port)
```

This function locks the driver used by the port `port` in memory for the rest of the emulator process lifetime. After this call, the driver behaves as one of Erlang's statically linked in drivers.

```
ErlDrvPort driver_create_port(ErlDrvPort port, ErlDrvTermData owner_pid, char* name,
                             ErlDrvData drv_data)
```

This function creates a new port executing the same driver code as the port creating the new port. A short description of the arguments:

`port` The port handle of the port (driver instance) creating the new port.

`owner_pid` The process id of the Erlang process which will be owner of the new port. This process will be linked to the new port. You usually want to use `driver_caller(port)` as `owner_pid`.

`name` The port name of the new port. You usually want to use the same port name as the driver name (`driver_name` [page 113] field of the `driver_entry` [page 110]).

`drv_data` The driver defined handle that will be passed in subsequent calls to driver call-backs. Note, that the driver start call-back [page 112] will not be called for this new driver instance. The driver defined handle is normally created in the driver start call-back [page 112] when a port is created via `[erlang:open_port/2]`.

The caller of `driver_create_port()` is allowed to manipulate the newly created port when `driver_create_port()` has returned. When port level locking [page 116] is used, the creating port is, however, only allowed to manipulate the newly created port until the current driver call-back that was called by the emulator returns.

Note:

When port level locking [page 116] is used, the creating port is only allowed to manipulate the newly created port until the current driver call-back returns.

```
int erl_drv_thread_create(char *name, ErlDrvTid *tid, void * (*func)(void *), void
    *arg, ErlDrvThreadOpts *opts)
```

Arguments:

`name` A string identifying the created thread. It will be used to identify the thread in planned future debug functionality.

`tid` A pointer to a thread identifier variable.

`func` A pointer to a function to execute in the created thread.

`arg` A pointer to argument to the `func` function.

`opts` A pointer to thread options to use or NULL.

This function creates a new thread. On success 0 is returned; otherwise, an `errno` value is returned to indicate the error. The newly created thread will begin executing in the function pointed to by `func`, and `func` will be passed `arg` as argument. When `erl_drv_thread_create()` returns the thread identifier of the newly created thread will be available in `*tid`. `opts` can be either a NULL pointer, or a pointer to an `ErlDrvThreadOpts` [page 122] structure. If `opts` is a NULL pointer, default options will be used; otherwise, the passed options will be used.

Warning:

You are not allowed to allocate the `ErlDrvThreadOpts` [page 122] structure by yourself. It has to be allocated and initialized by `erl_drv_thread_opts_create()` [page 137].

The created thread will terminate either when `func` returns or if `erl_drv_thread_exit()` [page 137] is called by the thread. The exit value of the thread is either returned from `func` or passed as argument to `erl_drv_thread_exit()` [page 137]. The driver creating the thread has the responsibility of joining the thread, via `erl_drv_thread_join()` [page 137], before the driver is unloaded. It is not possible to create “detached” threads, i.e., threads that don’t need to be joined.

Warning:

All created threads need to be joined by the driver before it is unloaded. If the driver fails to join all threads created before it is unloaded, the runtime system will most likely crash when the code of the driver is unloaded.

This function is thread-safe.

```
ErlDrvThreadOpts *erl_drv_thread_opts_create(char *name)
```

Arguments:

`name` A string identifying the created thread options. It will be used to identify the thread options in planned future debug functionality.

This function allocates and initialize a thread option structure. On failure NULL is returned. A thread option structure is used for passing options to `erl_drv_thread_create()` [page 136]. If the structure isn't modified before it is passed to `erl_drv_thread_create()` [page 136], the default values will be used.

Warning:

You are not allowed to allocate the `ErlDrvThreadOpts` [page 122] structure by yourself. It has to be allocated and initialized by `erl_drv_thread_opts_create()`.

This function is thread-safe.

```
void erl_drv_thread_opts_destroy(ErlDrvThreadOpts *opts)
```

Arguments:

`opts` A pointer to thread options to destroy.

This function destroys thread options previously created by `erl_drv_thread_opts_create()` [page 137].

This function is thread-safe.

```
void erl_drv_thread_exit(void *exit_value)
```

Arguments:

`exit_value` A pointer to an exit value or NULL.

This function terminates the calling thread with the exit value passed as argument. You are only allowed to terminate threads created with `erl_drv_thread_create()` [page 136]. The exit value can later be retrieved by another thread via `erl_drv_thread_join()` [page 137].

This function is thread-safe.

```
int erl_drv_thread_join(ErlDrvTid tid, void **exit_value)
```

Arguments:

`tid` The thread identifier of the thread to join.
`exit_value` A pointer to a pointer to an exit value, or NULL.

This function joins the calling thread with another thread, i.e., the calling thread is blocked until the thread identified by `tid` has terminated. On success 0 is returned; otherwise, an `errno` value is returned to indicate the error. A thread can only be joined once. The behavior of joining more than once is undefined, an emulator crash is likely. If `exit_value == NULL`, the exit value of the terminated thread will be ignored; otherwise, the exit value of the terminated thread will be stored at `*exit_value`. This function is thread-safe.

```
ErlDrvTid erl_drv_thread_self(void)
```

This function returns the thread identifier of the calling thread.
This function is thread-safe.

```
int erl_drv_equal_tids(ErlDrvTid tid1, ErlDrvTid tid2)
```

Arguments:

`tid1` A thread identifier.
`tid2` A thread identifier.

This function compares two thread identifiers for equality, and returns 0 if they aren't equal, and a value not equal to 0 if they are equal.

Note:

A Thread identifier may be reused very quickly after a thread has terminated. Therefore, if a thread corresponding to one of the involved thread identifiers has terminated since the thread identifier was saved, the result of `erl_drv_equal_tids()` might not give expected result.

This function is thread-safe.

```
ErlDrvMutex *erl_drv_mutex_create(char *name)
```

Arguments:

`name` A string identifying the created mutex. It will be used to identify the mutex in planned future debug functionality.

This function creates a mutex and returns a pointer to it. On failure NULL is returned. The driver creating the mutex has the responsibility of destroying it before the driver is unloaded.

This function is thread-safe.

```
void erl_drv_mutex_destroy(ErlDrvMutex *mtx)
```

Arguments:

`mtx` A pointer to a mutex to destroy.

This function destroys a mutex previously created by `erl_drv_mutex_create()` [page 138]. The mutex has to be in an unlocked state before being destroyed.

This function is thread-safe.

```
void erl_drv_mutex_lock(ErlDrvMutex *mtx)
```

Arguments:

`mtx` A pointer to a mutex to lock.

This function locks a mutex. The calling thread will be blocked until the mutex has been locked. A thread which currently has locked the mutex may *not* lock the same mutex again.

Warning:

If you leave a mutex locked in an emulator thread when you let the thread out of your control, you will *very likely* deadlock the whole emulator.

This function is thread-safe.

```
int erl_drv_mutex_trylock(ErlDrvMutex *mtx)
```

Arguments:

`mtx` A pointer to a mutex to try to lock.

This function tries to lock a mutex. If successful 0, is returned; otherwise, `EBUSY` is returned. A thread which currently has locked the mutex may *not* try to lock the same mutex again.

Warning:

If you leave a mutex locked in an emulator thread when you let the thread out of your control, you will *very likely* deadlock the whole emulator.

This function is thread-safe.

```
void erl_drv_mutex_unlock(ErlDrvMutex *mtx)
```

Arguments:

`mtx` A pointer to a mutex to unlock.

This function unlocks a mutex. The mutex currently has to be locked by the calling thread.

This function is thread-safe.

```
ErlDrvCond *erl_drv_cond_create(char *name)
```

Arguments:

`name` A string identifying the created condition variable. It will be used to identify the condition variable in planned future debug functionality.

This function creates a condition variable and returns a pointer to it. On failure `NULL` is returned. The driver creating the condition variable has the responsibility of destroying it before the driver is unloaded.

This function is thread-safe.

```
void erl_drv_cond_destroy(ErlDrvCond *cnd)
```

Arguments:

`cnd` A pointer to a condition variable to destroy.

This function destroys a condition variable previously created by `erl_drv_cond_create()` [page 139].

This function is thread-safe.

```
void erl_drv_cond_signal(ErlDrvCond *cnd)
```

Arguments:

`cnd` A pointer to a condition variable to signal on.

This function signals on a condition variable. That is, if other threads are waiting on the condition variable being signaled, *one* of them will be woken.

This function is thread-safe.

```
void erl_drv_cond_broadcast(ErlDrvCond *cnd)
```

Arguments:

`cnd` A pointer to a condition variable to broadcast on.

This function broadcasts on a condition variable. That is, if other threads are waiting on the condition variable being broadcasted on, *all* of them will be woken.

This function is thread-safe.

```
void erl_drv_cond_wait(ErlDrvCond *cnd, ErlDrvMutex *mtx)
```

Arguments:

`cnd` A pointer to a condition variable to wait on.

`mtx` A pointer to a mutex to unlock while waiting.

This function waits on a condition variable. The calling thread is blocked until another thread wakes it by signaling or broadcasting on the condition variable. Before the calling thread is blocked it unlocks the mutex passed as argument, and when the calling thread is woken it locks the same mutex before returning. That is, the mutex currently has to be locked by the calling thread when calling this function.

Note:

`erl_drv_cond_wait()` might return even though no-one has signaled or broadcasted on the condition variable. Code calling `erl_drv_cond_wait()` should always be prepared for `erl_drv_cond_wait()` returning even though the condition that the thread was waiting for hasn't occurred. That is, when returning from `erl_drv_cond_wait()` always check if the condition has occurred, and if not call `erl_drv_cond_wait()` again.

This function is thread-safe.

```
ErlDrvRWLock *erl_drv_rwlock_create(char *name)
```

Arguments:

`name` A string identifying the created rwlock. It will be used to identify the rwlock in planned future debug functionality.

This function creates an rwlock and returns a pointer to it. On failure NULL is returned. The driver creating the rwlock has the responsibility of destroying it before the driver is unloaded.

This function is thread-safe.

```
void erl_drv_rwlock_destroy(ErlDrvRWLock *rwlock)
```

Arguments:

`rwlock` A pointer to an rwlock to destroy.

This function destroys an rwlock previously created by `erl_drv_rwlock_create()` [page 141]. The rwlock has to be in an unlocked state before being destroyed.

This function is thread-safe.

```
void erl_drv_rwlock_rlock(ErlDrvRWLock *rwlock)
```

Arguments:

`rwlock` A pointer to an rwlock to read lock.

This function read locks an rwlock. The calling thread will be blocked until the rwlock has been read locked. A thread which currently has read or read/write locked the rwlock may *not* lock the same rwlock again.

Warning:

If you leave an rwlock locked in an emulator thread when you let the thread out of your control, you will *very likely* deadlock the whole emulator.

This function is thread-safe.

```
int erl_drv_rwlock_tryrlock(ErlDrvRWLock *rwlock)
```

Arguments:

`rwlock` A pointer to an rwlock to try to read lock.

This function tries to read lock an rwlock. If successful 0, is returned; otherwise, EBUSY is returned. A thread which currently has read or read/write locked the rwlock may *not* try to lock the same rwlock again.

Warning:

If you leave an rwlock locked in an emulator thread when you let the thread out of your control, you will *very likely* deadlock the whole emulator.

This function is thread-safe.

```
void erl_drv_rwlock_runlock(ErlDrvRWLock *rwlock)
```

Arguments:

`rwlock` A pointer to an rwlock to read unlock.

This function read unlocks an rwlock. The rwlock currently has to be read locked by the calling thread.

This function is thread-safe.

```
void erl_drv_rwlock_rwlock(ErlDrvRWLock *rwlock)
```

Arguments:

`rwlock` A pointer to an rwlock to read/write lock.

This function read/write locks an rwlock. The calling thread will be blocked until the rwlock has been read/write locked. A thread which currently has read or read/write locked the rwlock may *not* lock the same rwlock again.

Warning:

If you leave an rwlock locked in an emulator thread when you let the thread out of your control, you will *very likely* deadlock the whole emulator.

This function is thread-safe.

```
int erl_drv_rwlock_tryrwlock(ErlDrvRWLock *rwlock)
```

Arguments:

`rwlock` A pointer to an rwlock to try to read/write lock.

This function tries to read/write lock an rwlock. If successful 0, is returned; otherwise, `EBUSY` is returned. A thread which currently has read or read/write locked the rwlock may *not* try to lock the same rwlock again.

Warning:

If you leave an rwlock locked in an emulator thread when you let the thread out of your control, you will *very likely* deadlock the whole emulator.

This function is thread-safe.

```
void erl_drv_rwlock_rwlock(ErlDrvRWLock *rwlock)
```

Arguments:

`rwlock` A pointer to an rwlock to read/write unlock.

This function read/write unlocks an rwlock. The rwlock currently has to be read/write locked by the calling thread.

This function is thread-safe.

```
int erl_drv_tsd_key_create(char *name, ErlDrvTSDKey *key)
```

Arguments:

`name` A string identifying the created key. It will be used to identify the key in planned future debug functionality.

`key` A pointer to a thread specific data key variable.

This function creates a thread specific data key. On success 0 is returned; otherwise, an `errno` value is returned to indicate the error. The driver creating the key has the responsibility of destroying it before the driver is unloaded.

This function is thread-safe.

```
void erl_drv_tsd_key_destroy(ErlDrvTSDKey key)
```

Arguments:

`key` A thread specific data key to destroy.

This function destroys a thread specific data key previously created by `erl_drv_tsd_key_create()` [page 143]. All thread specific data using this key in all threads have to be cleared (see `erl_drv_tsd_set()` [page 144]) prior to the call to `erl_drv_tsd_key_destroy()`.

Warning:

A destroyed key is very likely to be reused soon. Therefore, if you fail to clear the thread specific data using this key in a thread prior to destroying the key, you will *very likely* get unexpected errors in other parts of the system.

This function is thread-safe.

```
void erl_drv_tsd_set(ErlDrvTSDKey key, void *data)
```

Arguments:

`key` A thread specific data key.

`data` A pointer to data to associate with `key` in calling thread.

This function sets thread specific data associated with `key` for the calling thread. You are only allowed to set thread specific data for threads while they are fully under your control. For example, if you set thread specific data in a thread calling a driver call-back function, it has to be cleared, i.e. set to `NULL`, before returning from the driver call-back function.

Warning:

If you fail to clear thread specific data in an emulator thread before letting it out of your control, you might not ever be able to clear this data with later unexpected errors in other parts of the system as a result.

This function is thread-safe.

```
void *erl_drv_tsd_get(ErlDrvTSDKey key)
```

Arguments:

`key` A thread specific data key.

This function returns the thread specific data associated with `key` for the calling thread. If no data has been associated with `key` for the calling thread, `NULL` is returned.

This function is thread-safe.

```
int erl_drv_putenv(char *key, char *value)
```

Arguments:

`key` A null terminated string containing the name of the environment variable.

`value` A null terminated string containing the new value of the environment variable.

This function sets the value of an environment variable. It returns 0 on success, and a value $\neq 0$ on failure.

Note:

The result of passing the empty string (“”) as a value is platform dependent. On some platforms the value of the variable is set to the empty string, on others, the environment variable is removed.

Warning:

Do *not* use libc’s `putenv` or similar C library interfaces from a driver.

This function is thread-safe.

```
int erl_drv_getenv(char *key, char *value, size_t *value_size)
```

Arguments:

`key` A null terminated string containing the name of the environment variable.

`value` A pointer to an output buffer.

`value_size` A pointer to an integer. The integer is both used for passing input and output sizes (see below).

This function retrieves the value of an environment variable. When called, `*value_size` should contain the size of the `value` buffer. On success 0 is returned, the value of the environment variable has been written to the `value` buffer, and `*value_size` contains the string length (excluding the terminating null character) of the value written to the `value` buffer. On failure, i.e., no such environment variable was found, a value less than 0 is returned. When the size of the `value` buffer is too small, a value greater than 0 is returned and `*value_size` has been set to the buffer size needed.

Warning:

Do *not* use libc’s `getenv` or similar C library interfaces from a driver.

This function is thread-safe.

SEE ALSO

`driver_entry(3)` [page 110], `[erl_dll(3)]`, `[erlang(3)]`

An Alternative Distribution Driver (ERTS User’s Guide Ch. 3)

List of Figures

1.1 Summary of EPMD requests. 60

List of Tables

1.1	Literals in the MatchCondition/MatchBody parts of a match_spec	5
1.2	tty text editing	35
1.3	51
1.4	51
1.5	52
1.6	52
1.7	52
1.8	52
1.9	53
1.10	53
1.11	53
1.12	54
1.13	54
1.14	54
1.15	54
1.16	55
1.17	55
1.18	55
1.19	55
1.20	56
1.21	56
1.22	56
1.23	57
1.24	57
1.25	57
1.26	58
1.27	58
1.28	59
1.29	61
1.30	ALIVE2_REQ (120)	61

1.31 ALIVE2_RESP (121)	61
1.32 PORT_PLEASE2_REQ (122)	62
1.33 PORT2_RESP (119) response indicating error, Result > 0.	62
1.34 PORT2_RESP when Result = 0.	62
1.35 NAMES_REQ (110)	62
1.36 NAMES_RESP	63
1.37 DUMP_REQ	63
1.38 DUMP_RESP	63
1.39 KILL_REQ	64
1.40 KILL_RESP	64
1.41 STOP_REQ	64
1.42 STOP_RESP	64
1.43 STOP_NOTOK_RESP	65
1.44	65

Index of Modules and Functions

Modules are typed in *this way*.
Functions are typed in *this way*.

char *driver_name
 driver_entry, 113

char* erl_errno_id/1
 erl_driver, 131

driver_entry
 char *driver_name, 113
 int call/7, 114
 int control/6, 113
 int driver_flags, 115
 int extended_marker, 114
 int init/1, 112
 int major_version, 115
 int minor_version, 115
 int start/2, 112
 void *handle, 113
 void *handle2, 115
 void event/3, 114
 void finish/1, 113
 void output/3, 112
 void outputv/2, 114
 void process_exit/2, 115
 void ready_async/2, 114
 void ready_input/2, 112
 void ready_output/2, 112
 void stop/1, 112
 void timeout/1, 114

epmd (Command)
 epmd, 78

epmd
 epmd (Command), 78

erl (Command)
 erl, 79

erl
 erl (Command), 79

erl_driver
 char* erl_errno_id/1, 131

ErlDrvBinary, 120
ErlDrvBinary* driver_alloc_binary/1,
 126
ErlDrvBinary* driver_realloc_binary/2,
 126
ErlDrvCond, 122
ErlDrvCond *erl_drv_cond_create/1,
 139
ErlDrvData, 121
ErlDrvMonitor, 121
ErlDrvMutex, 122
ErlDrvMutex *erl_drv_mutex_create/1,
 138
ErlDrvNowData, 121
ErlDrvPDL, 122
ErlDrvPDL driver_pdl_create/1, 129
ErlDrvPort driver_create_port/4, 135
ErlDrvRWLock, 123
ErlDrvRWLock *erl_drv_rwlock_create/1,
 141
ErlDrvSysInfo, 119
ErlDrvTermData driver_caller/1, 131
ErlDrvTermData driver_connected/1,
 131
ErlDrvTermData driver_get_monitored_process/2,
 130
ErlDrvTermData driver_mk_atom/1, 134
ErlDrvTermData driver_mk_port/1, 134
ErlDrvThreadOpts, 122
ErlDrvThreadOpts *erl_drv_thread_opts_create/1,
 137
ErlDrvTid, 122
ErlDrvTid erl_drv_thread_self/1, 138
ErlDrvTSDKey, 123
ErlIOVec, 121
int driver_async_cancel/1, 135
int driver_cancel_timer/1, 125
int driver_compare_monitors/2, 130
int driver_demonitor_process/2, 130
int driver_deq/2, 127
int driver_enq/3, 127

int driver_enq_bin/4, 128
 int driver_enqv/3, 128
 int driver_failure/2, 131
 int driver_failure_atom/2, 131
 int driver_failure_eof/1, 131
 int driver_failure_posix/2, 131
 int driver_get_now/1, 125
 int driver_lock_driver/1, 135
 int driver_monitor_process/3, 129
 int driver_output/3, 123
 int driver_output2/5, 123
 int driver_output_binary/6, 124
 int driver_output_term/3, 132
 int driver_outputv/5, 124
 int driver_pushq/3, 127
 int driver_pushq_bin/4, 128
 int driver_pushqv/3, 129
 int driver_read_timer/2, 125
 int driver_select/4, 125
 int driver_send_term/4, 134
 int driver_set_timer/2, 124
 int driver_sizeq/1, 128
 int driver_vec_to_buf/3, 124
 int erl_drv_equal_tids/2, 138
 int erl_drv_getenv/3, 145
 int erl_drv_mutex_trylock/1, 139
 int erl_drv_putenv/2, 144
 int erl_drv_rwlock_tryrlock/1, 142
 int erl_drv_rwlock_tryrwlock/1, 143
 int erl_drv_thread_create/3, 136
 int erl_drv_thread_join/2, 137
 int erl_drv_tsd_key_create/2, 143
 int remove_driver_entry/1, 130
 long driver_async/3, 134
 long driver_binary_dec_refc/1, 127
 long driver_binary_get_refc/1, 126
 long driver_binary_inc_refc/1, 127
 long driver_pdl_dec_refc/1, 129
 long driver_pdl_get_refc/1, 129
 long driver_pdl_inc_refc/1, 129
 SysIOVec, 121
 SysIOVec* driver_peekq/2, 128
 void *driver_alloc/1, 125
 void *driver_realloc/2, 126
 void *erl_drv_tsd_get/1, 144
 void add_driver_entry/1, 130
 void driver_free/1, 126
 void driver_free_binary/1, 126
 void driver_pdl_lock/1, 129
 void driver_pdl_unlock/1, 129
 void driver_system_info/2, 123
 void erl_drv_cond_broadcast/1, 140
 void erl_drv_cond_destroy/1, 140
 void erl_drv_cond_signal/1, 140
 void erl_drv_cond_wait/2, 140
 void erl_drv_mutex_destroy/1, 138
 void erl_drv_mutex_lock/1, 139
 void erl_drv_mutex_unlock/1, 139
 void erl_drv_rwlock_destroy/1, 141
 void erl_drv_rwlock_rlock/1, 141
 void erl_drv_rwlock_runlock/1, 142
 void erl_drv_rwlock_rwlock/1, 142
 void erl_drv_rwlock_runlock/1, 143
 void erl_drv_thread_exit/1, 137
 void erl_drv_thread_opts_destroy/1, 137
 void erl_drv_tsd_key_destroy/1, 144
 void erl_drv_tsd_set/2, 144
 void set_busy_port/2, 131
 void set_port_control_flags/2, 131
 erl_mem_info_get/1 (C function)
 erl_set_memory_block, 103
 erl_memory_show/1 (C function)
 erl_set_memory_block, 103
erl_set_memory_block
 erl_mem_info_get/1 (C function), 103
 erl_memory_show/1 (C function), 103
 erl_set_memory_block/6 (C function), 102
 erl_set_memory_block/6 (C function)
 erl_set_memory_block, 102
 erlc (Command)
 erlc, 86
erlc
 erlc (Command), 86
 ErlDrvBinary
 erl_driver, 120
 ErlDrvBinary* driver_alloc_binary/1
 erl_driver, 126
 ErlDrvBinary* driver_realloc_binary/2
 erl_driver, 126
 ErlDrvCond
 erl_driver, 122
 ErlDrvCond *erl_drv_cond_create/1
 erl_driver, 139
 ErlDrvData
 erl_driver, 121
 ErlDrvMonitor
 erl_driver, 121

ErlDrvMutex	erlsrv (Command), 91, 92
<i>erl_driver</i> , 122	
ErlDrvMutex *erl_drv_mutex_create/1	escript (Command)
<i>erl_driver</i> , 138	<i>escript</i> , 94
ErlDrvNowData	<i>escript</i>
<i>erl_driver</i> , 121	escript (Command), 94
ErlDrvPDL	script-name (Command), 94
<i>erl_driver</i> , 122	
ErlDrvPDL driver_pdl_create/1	int call/7
<i>erl_driver</i> , 129	<i>driver_entry</i> , 114
ErlDrvPort driver_create_port/4	int control/6
<i>erl_driver</i> , 135	<i>driver_entry</i> , 113
ErlDrvRWLock	int driver_async_cancel/1
<i>erl_driver</i> , 123	<i>erl_driver</i> , 135
ErlDrvRWLock *erl_drv_rwlock_create/1	int driver_cancel_timer/1
<i>erl_driver</i> , 141	<i>erl_driver</i> , 125
ErlDrvSysInfo	int driver_compare_monitors/2
<i>erl_driver</i> , 119	<i>erl_driver</i> , 130
ErlDrvTermData driver_caller/1	int driver_demonitor_process/2
<i>erl_driver</i> , 131	<i>erl_driver</i> , 130
ErlDrvTermData driver_connected/1	int driver_deq/2
<i>erl_driver</i> , 131	<i>erl_driver</i> , 127
ErlDrvTermData driver_get_monitored_process/2	int driver_enq/3
<i>erl_driver</i> , 130	<i>erl_driver</i> , 127
ErlDrvTermData driver_mk_atom/1	int driver_enq_bin/4
<i>erl_driver</i> , 134	<i>erl_driver</i> , 128
ErlDrvTermData driver_mk_port/1	int driver_enqv/3
<i>erl_driver</i> , 134	<i>erl_driver</i> , 128
ErlDrvThreadOpts	int driver_failure/2
<i>erl_driver</i> , 122	<i>erl_driver</i> , 131
ErlDrvThreadOpts *erl_drv_thread_opts_create/1	int driver_failure_atom/2
<i>erl_driver</i> , 137	<i>erl_driver</i> , 131
ErlDrvTid	int driver_failure_eof/1
<i>erl_driver</i> , 122	<i>erl_driver</i> , 131
ErlDrvTid erl_drv_thread_self/1	int driver_failure_posix/2
<i>erl_driver</i> , 138	<i>erl_driver</i> , 131
ErlDrvTSDKey	int driver_flags
<i>erl_driver</i> , 123	<i>driver_entry</i> , 115
ErlIOVec	int driver_get_now/1
<i>erl_driver</i> , 121	<i>erl_driver</i> , 125
erlsrv (Command)	int driver_lock_driver/1
<i>erlsrv</i> , 91, 92	<i>erl_driver</i> , 135
<i>erlsrv</i>	int driver_monitor_process/3
	<i>erl_driver</i> , 129

int driver_output/3
 erl_driver, 123
 int driver_output2/5
 erl_driver, 123
 int driver_output_binary/6
 erl_driver, 124
 int driver_output_term/3
 erl_driver, 132
 int driver_outputv/5
 erl_driver, 124
 int driver_pushq/3
 erl_driver, 127
 int driver_pushq_bin/4
 erl_driver, 128
 int driver_pushqv/3
 erl_driver, 129
 int driver_read_timer/2
 erl_driver, 125
 int driver_select/4
 erl_driver, 125
 int driver_send_term/4
 erl_driver, 134
 int driver_set_timer/2
 erl_driver, 124
 int driver_sizeq/1
 erl_driver, 128
 int driver_vec_to_buf/3
 erl_driver, 124
 int erl_drv_equal_tids/2
 erl_driver, 138
 int erl_drv_getenv/3
 erl_driver, 145
 int erl_drv_mutex_trylock/1
 erl_driver, 139
 int erl_drv_putenv/2
 erl_driver, 144
 int erl_drv_rwlock_trylock/1
 erl_driver, 142
 int erl_drv_rwlock_tryrlock/1
 erl_driver, 143
 int erl_drv_thread_create/3
 erl_driver, 136
 int erl_drv_thread_join/2
 erl_driver, 137
 int erl_drv_tsd_key_create/2
 erl_driver, 143
 int extended_marker
 driver_entry, 114
 int init/1
 driver_entry, 112
 int major_version
 driver_entry, 115
 int minor_version
 driver_entry, 115
 int remove_driver_entry/1
 erl_driver, 130
 int start/2
 driver_entry, 112
 long driver_async/3
 erl_driver, 134
 long driver_binary_dec_refc/1
 erl_driver, 127
 long driver_binary_get_refc/1
 erl_driver, 126
 long driver_binary_inc_refc/1
 erl_driver, 127
 long driver_pdl_dec_refc/1
 erl_driver, 129
 long driver_pdl_get_refc/1
 erl_driver, 129
 long driver_pdl_inc_refc/1
 erl_driver, 129
 run_erl (Command)
 run_erl, 96
run_erl
 run_erl (Command), 96
 script-name (Command)
 escript, 94
 start (Command)
 start, 98
start
 start (Command), 98
 start_erl (Command)
 start_erl, 99

<i>start_erl</i>	
<i>start_erl</i> (Command), 99	
<i>SysIOVec</i>	
<i>erl_driver</i> , 121	
<i>SysIOVec* driver_peekq/2</i>	
<i>erl_driver</i> , 128	
<i>void *driver_alloc/1</i>	
<i>erl_driver</i> , 125	
<i>void *driver_realloc/2</i>	
<i>erl_driver</i> , 126	
<i>void *erl_drv_tsd_get/1</i>	
<i>erl_driver</i> , 144	
<i>void *handle</i>	
<i>driver_entry</i> , 113	
<i>void *handle2</i>	
<i>driver_entry</i> , 115	
<i>void add_driver_entry/1</i>	
<i>erl_driver</i> , 130	
<i>void driver_free/1</i>	
<i>erl_driver</i> , 126	
<i>void driver_free_binary/1</i>	
<i>erl_driver</i> , 126	
<i>void driver_pdl_lock/1</i>	
<i>erl_driver</i> , 129	
<i>void driver_pdl_unlock/1</i>	
<i>erl_driver</i> , 129	
<i>void driver_system_info/2</i>	
<i>erl_driver</i> , 123	
<i>void erl_drv_cond_broadcast/1</i>	
<i>erl_driver</i> , 140	
<i>void erl_drv_cond_destroy/1</i>	
<i>erl_driver</i> , 140	
<i>void erl_drv_cond_signal/1</i>	
<i>erl_driver</i> , 140	
<i>void erl_drv_cond_wait/2</i>	
<i>erl_driver</i> , 140	
<i>void erl_drv_mutex_destroy/1</i>	
<i>erl_driver</i> , 138	
<i>void erl_drv_mutex_lock/1</i>	
<i>erl_driver</i> , 139	
<i>void erl_drv_mutex_unlock/1</i>	
<i>erl_driver</i> , 139	
<i>void erl_drv_rwlock_destroy/1</i>	
<i>erl_driver</i> , 141	
<i>void erl_drv_rwlock_rlock/1</i>	
<i>erl_driver</i> , 141	
<i>void erl_drv_rwlock_runlock/1</i>	
<i>erl_driver</i> , 142	
<i>void erl_drv_rwlock_rwlock/1</i>	
<i>erl_driver</i> , 142	
<i>void erl_drv_rwlock_rwunlock/1</i>	
<i>erl_driver</i> , 143	
<i>void erl_drv_thread_exit/1</i>	
<i>erl_driver</i> , 137	
<i>void erl_drv_thread_opts_destroy/1</i>	
<i>erl_driver</i> , 137	
<i>void erl_drv_tsd_key_destroy/1</i>	
<i>erl_driver</i> , 144	
<i>void erl_drv_tsd_set/2</i>	
<i>erl_driver</i> , 144	
<i>void event/3</i>	
<i>driver_entry</i> , 114	
<i>void finish/1</i>	
<i>driver_entry</i> , 113	
<i>void output/3</i>	
<i>driver_entry</i> , 112	
<i>void outputv/2</i>	
<i>driver_entry</i> , 114	
<i>void process_exit/2</i>	
<i>driver_entry</i> , 115	
<i>void ready_async/2</i>	
<i>driver_entry</i> , 114	
<i>void ready_input/2</i>	
<i>driver_entry</i> , 112	
<i>void ready_output/2</i>	
<i>driver_entry</i> , 112	
<i>void set_busy_port/2</i>	
<i>erl_driver</i> , 131	
<i>void set_port_control_flags/2</i>	
<i>erl_driver</i> , 131	
<i>void stop/1</i>	
<i>driver_entry</i> , 112	
<i>void timeout/1</i>	
<i>driver_entry</i> , 114	

