

# **Kernel Application (KERNEL)**

**version 2.2**

**Magnus Fröberg**

**1997-11-17**

Typeset in  $\text{\LaTeX}$  from SGML source using the DOCBUILDER 3.0 Document System.



# Contents

<b>1</b>	<b>KERNEL Reference Manual</b>	<b>1</b>
1.1	kernel (Application)	24
1.2	application (Module)	28
1.3	auth (Module)	35
1.4	code (Module)	38
1.5	disk_log (Module)	46
1.6	erl_boot_server (Module)	54
1.7	erl_dll (Module)	56
1.8	erl_prim_loader (Module)	59
1.9	erlang (Module)	62
1.10	error_handler (Module)	87
1.11	error_logger (Module)	89
1.12	file (Module)	94
1.13	gen_tcp (Module)	107
1.14	gen_udp (Module)	111
1.15	global (Module)	113
1.16	global_group (Module)	117
1.17	heart (Module)	121
1.18	inet (Module)	123
1.19	init (Module)	130
1.20	net_adm (Module)	134
1.21	net_kernel (Module)	136
1.22	os (Module)	137
1.23	pg2 (Module)	139
1.24	rpc (Module)	141
1.25	seq_trace (Module)	145
1.26	user (Module)	153
1.27	wrap_log_reader (Module)	154
1.28	app (File)	156
1.29	config (File)	158



# KERNEL Reference Manual

## Short Summaries

- Application **kernel** [page 24] – The Kernel Application
- Erlang Module **application** [page 28] – Functions for controlling applications
- Erlang Module **auth** [page 35] – The Erlang Network Authentication Server
- Erlang Module **code** [page 38] – Erlang Code Server
- Erlang Module **disk\_log** [page 46] – A disk based term logging facility
- Erlang Module **erl\_boot\_server** [page 54] – Boot Server for Other Erlang Machines
- Erlang Module **erl\_ddll** [page 56] – Dynamic Driver Loader and Linker
- Erlang Module **erl\_prim\_loader** [page 59] – The Low Level Erlang Loader.
- Erlang Module **erlang** [page 62] – The Erlang BIFs
- Erlang Module **error\_handler** [page 87] – Default System Error Handler
- Erlang Module **error\_logger** [page 89] – The Erlang Error Logger
- Erlang Module **file** [page 94] – File Interface Module
- Erlang Module **gen\_tcp** [page 107] – Interface to TCP/IP sockets
- Erlang Module **gen\_udp** [page 111] – Interface to UDP.
- Erlang Module **global** [page 113] – A Global Name Registration Facility
- Erlang Module **global\_group** [page 117] – Grouping Nodes to Global Name Registration Groups
- Erlang Module **heart** [page 121] – Heartbeat Monitoring of an Erlang System.
- Erlang Module **inet** [page 123] – Access to TCP/IP protocols.
- Erlang Module **init** [page 130] – Called at System Start
- Erlang Module **net\_adm** [page 134] – Various Erlang Net Administration Routines
- Erlang Module **net\_kernel** [page 136] – Erlang Networking Kernel
- Erlang Module **os** [page 137] – Operating System Specific Functions
- Erlang Module **pg2** [page 139] – Distributed Named Process Groups
- Erlang Module **rpc** [page 141] – Remote Procedure Call Services
- Erlang Module **seq\_trace** [page 145] – Sequential Tracing of Messages.
- Erlang Module **user** [page 153] – Standard I/O Server
- Erlang Module **wrap\_log\_reader** [page 154] – A function to read internally formatted wrap disk logs
- File **app** [page 156] – Application resource file
- File **config** [page 158] – Configuration file

## kernel

No functions are exported

## application

The following functions are exported:

- `get_all_env()`  
[page 28] Gets all configuration parameters for the application.
- `get_all_env(Application) -> Env`  
[page 28] Gets all configuration parameters for the application.
- `get_all_key()`  
[page 28] Gets all the resource file keys for the application.
- `get_all_key(Application) -> {ok, Keys} | undefined`  
[page 28] Gets all the resource file keys for the application.
- `get_application()`  
[page 29] Gets the name of an application.
- `get_application(Pid | Module) -> {ok, Application} | undefined`  
[page 29] Gets the name of an application.
- `get_env(Key)`  
[page 29] Gets the value of a configuration parameter.
- `get_env(Application, Key) -> {ok, Value} | undefined`  
[page 29] Gets the value of a configuration parameter.
- `get_key(Key)`  
[page 29] Gets the value of a configuration parameter.
- `get_key(Application, Key) -> {ok, Value} | undefined`  
[page 29] Gets the value of a configuration parameter.
- `load(Application)`  
[page 29] Loads an application.
- `load(Application, DistNodes) -> ok | {error, Reason}`  
[page 29] Loads an application.
- `loaded_applications() -> [{Name, Description, Version}]`  
[page 30] Gets the currently loaded applications
- `permit(Application, Bool) -> ok | {error, Reason}`  
[page 30] Changes an application's permission to run on the node
- `start(Application)`  
[page 31] Starts an application.
- `start(Application, Type) -> ok | {error, Reason}`  
[page 31] Starts an application.
- `start_type() -> normal | local | {takeover, node()} | {failover, node()}`  
[page 31] Check the type of start of an application.
- `stop(Application) -> ok`  
[page 31] Stops a running application

- `takeover(Application, Type) -> {ok, Pid} | {error, Reason}`  
[page 32] Moves a distributed application to the current node
- `which_applications() -> [{Name, Description, Version}]`  
[page 32] Gets the currently running applications
- `Module:config_change(Changed, New, Removed) -> ok`  
[page 32] Informs an application of changed configuration parameters
- `Module:start(Type, ModuleStartArgs) -> {ok, Pid} | {ok, Pid, State} | {error, Reason}`  
[page 33] Starts an application
- `Module:start_phase(Phase, Type, PhaseStartArgs) -> ok | {error, Reason}`  
[page 33] Starts an application in the Phase
- `Module:stop(State) -> void()`  
[page 34] Called when the application has stopped

## auth

The following functions are exported:

- `start()`  
[page 36]
- `stop()`  
[page 36]
- `open(Name)`  
[page 36]
- `is_auth(Node)`  
[page 36]
- `exists(Node)`  
[page 36]
- `cookie()`  
[page 37]
- `node_cookie(Node, Cookie)`  
[page 37]
- `node_cookie([Node, Cookie])`  
[page 37]
- `cookie([Cookie])`  
[page 37]

## code

The following functions are exported:

- `start() -> {ok, Pid} | {error, What}`  
[page 38] Starts the code server.

- `start(Flags) -> {ok, Pid} | {error, What}`  
[page 38] Starts the code server.
- `start_link() -> {ok, Pid} | {error, What}`  
[page 38] Starts and links to the code server.
- `start_link(Flags) -> {ok, Pid} | {error, What}`  
[page 38] Starts and links to the code server.
- `set_path(DirList) -> true | {error, What}`  
[page 39] Sets the code server search path.
- `get_path() -> Path`  
[page 39] Returns the current path of the code server.
- `add_path(Dir) -> true | {error, What}`  
[page 39] Add a directory to the end of path.
- `add_pathz(Dir) -> true | {error, What}`  
[page 39] Add a directory to the end of path.
- `add_patha(Dir) -> true | {error, What}`  
[page 39] Adds a directory to the beginning of path.
- `add_paths(DirList) -> ok`  
[page 39] Adds directories to the end of path.
- `add_pathsz(DirList) -> ok`  
[page 39] Adds directories to the end of path.
- `add_pathsa(DirList) -> ok`  
[page 40] Adds directories to the beginning of path.
- `del_path(NameDir) -> true | false | {error, What}`  
[page 40] Deletes a directory from the path.
- `replace_path(Name, Dir) -> true | {error, What}`  
[page 40] Replaces a directory with another in the path.
- `load_file(Module) -> {module, Module} | {error, What}`  
[page 40] Loads a module (residing in File).
- `load_abs(File) -> {module, Module} | {error, What}`  
[page 40] Loads a module (residing in File).
- `ensure_loaded(Module) -> {module, Module} | {error, What} | {interpret, Module}`  
[page 41] Tries to ensure that a module is loaded.
- `delete(Module) -> true | false`  
[page 41] Deletes the code in Module.
- `purge(Module) -> true | false`  
[page 41] Purges the code in Module.
- `soft_purge(Module) -> true | false`  
[page 41] Purges the code in Module if no process uses it.
- `is_loaded(Module) -> {file, Loaded} | false`  
[page 41] Tests if Module is loaded.
- `all_loaded() -> [LoadMod]`  
[page 42] Gets all loaded modules.
- `load_binary(Module, File, Binary) -> {module, Module} | {error, What}`  
[page 42] Loads object code as a binary.

- `stop()` -> `stopped`  
[page 42] Stops the code server.
- `root_dir()` -> `RootDir`  
[page 42] Returns the root directory of the Erlang system.
- `lib_dir()` -> `LibDir`  
[page 42] Returns the library directory.
- `lib_dir(Name)` -> `LibDir` | `{error, What}`  
[page 42] Returns the directory for name.
- `compiler_dir()` -> `CompDir`  
[page 43] Returns the compiler directory.
- `priv_dir(Name)` -> `PrivDir` | `{error, What}`  
[page 43] Returns the priv directory for name.
- `get_object_code(Module)` -> `{Module, Bin, AbsFileName}` | `error`  
[page 43] Gets the object code for a module.
- `objfile_extension()` -> `Ext`  
[page 43] Returns the object code file extension.
- `stick_dir(Dir)` -> `ok` | `{error, term()}`  
[page 43] Marks a directory as 'sticky'.
- `unstick_dir(Dir)` -> `ok` | `{error, term()}`  
[page 44] Marks a directory as 'non-sticky'.
- `which(Module)` -> `WhichFile`  
[page 44] Returns the directory to a module.
- `clash()` -> `ok`  
[page 44] Searches for modules with identical names.
- `interpret(Module)` -> `{module, Module}` | `{error, What}`  
[page 44] Marks a module as being interpreted.
- `interpret_binary(Module, File, Binary)` -> `{module, Module}` | `{error, What}`  
[page 44] Loads an interpreted module into the interpreter.
- `delete_interpret(Module)` -> `ok` | `{error, What}`  
[page 44] Do not interpret a module.
- `interpreted()` -> `Modules`  
[page 45] Returns all interpreted modules.
- `interpreted(Module)` -> `true` | `false`  
[page 45]

## disk.log

The following functions are exported:

- `alog(Log, Term)` -> `ok` | `{error, Reason}`  
[page 46] Asynchronously logs an item
- `balog(Log, Bytes)` -> `ok` | `{error, Reason}`  
[page 46] Asynchronously logs an item
- `alog_terms(Log, TermList)` -> `ok` | `{error, Reason}`  
[page 47] Asynchronously logs several items

- `balog_terms(Log, BytesList) -> ok | {error, Reason}`  
[page 47] Asynchronously logs several items
- `block(Log)`  
[page 47] Blocks a log
- `block(Log, QueueLogRecords) -> ok | {error, Reason}`  
[page 47] Blocks a log
- `change_size(Log, Size) -> ok | {error, Reason}`  
[page 47] Changes the size of an open log
- `chunk(Log, Continuation)`  
[page 48] Reads a chunk of objects written to the log.
- `chunk(Log, Continuation, N) -> {Continuation2, Terms} | {Continuation2, Terms, Badbytes} | eof | {error, Reason}`  
[page 48] Reads a chunk of objects written to the log.
- `chunk_step(Log, Continuation, Step) -> {ok, Continuation2} | {error, Reason}`  
[page 49] Steps
- `close(Log) -> ok | {error, Reason}`  
[page 49] Closes a log
- `inc_wrap_log(Log) -> ok | {error, Reason}`  
[page 49] Change to next wrap log
- `log(Log, Term) -> ok | {error, Reason}`  
[page 49] Logs an item
- `blog(Log, Bytes) -> ok | {error, Reason}`  
[page 49] Logs an item
- `log_terms(Log, TermList) -> ok | {error, Reason}`  
[page 50] Logs several items
- `blog_terms(Log, BytesList) -> ok | {error, Reason}`  
[page 50] Logs several items
- `open(ArgL) -> OpenRet`  
[page 50] Opens a log file
- `reopen(Log, File)`  
[page 52] Reopens the log file, and saves the old log
- `reopen(Log, File, Head)`  
[page 52] Reopens the log file, and saves the old log
- `breopen(Log, File, BHead) -> ok | {error, Reason}`  
[page 52] Reopens the log file, and saves the old log
- `sync(Log) -> ok | {error, Reason}`  
[page 52] Flushes contents of the log to the disk
- `truncate(Log)`  
[page 52] Truncates a log
- `truncate(Log, Head)`  
[page 52] Truncates a log
- `btruncate(Log, BHead) -> ok | {error, Reason}`  
[page 52] Truncates a log
- `unblock(Log) -> ok | {error, Reason}`  
[page 52] Unblocks a log

## erl\_boot\_server

The following functions are exported:

- `start(Slaves) -> {ok, Pid} | {error, What}`  
[page 54] Starts the boot server.
- `start_link(Slaves) -> {ok, Pid} | {error, What}`  
[page 54] Starts the boot server and links the caller.
- `add_slave(Slave) -> ok | {error, What}`  
[page 54] Adds a slave to the list of allowed slaves.
- `delete_slave(Slave) -> ok | {error, What}`  
[page 55] Deletes a slave from the list of allowed slaves.
- `which_slaves() -> Slaves`  
[page 55] Returns the current list of allowed slave hosts.

## erl\_ddll

The following functions are exported:

- `start() -> {ok, Pid} | {error, Reason}`  
[page 56] Starts the server.
- `start_link() -> {ok, Pid} | {error, Reason}`  
[page 56] Starts the server and links it to the calling process.
- `stop() -> ok`  
[page 56] Stops the server.
- `load_driver(Path, Name) -> ok | {error, ErrorDescriptor}`  
[page 56] Loads a driver.
- `unload_driver(Name) -> ok | {error, ErrorDescriptor}`  
[page 56] Unloads a driver.
- `loaded_drivers() -> {ok, DriverList}`  
[page 57] Lists loaded drivers.
- `format_error(ErrorDescriptor) -> string()`  
[page 57] Formats an error descriptor

## erl\_prim\_loader

The following functions are exported:

- `start(Id, Loader, Hosts) -> {ok, Pid} | {error, What}`  
[page 59] Starts the Erlang low level loader.
- `get_file(File) -> {ok, Bin, FullName} | error`  
[page 59] Gets a file.
- `get_path() -> {ok, Path}`  
[page 60] Gets the path set in the loader.
- `set_path(Path) -> ok`  
[page 60] Sets the path of the loader.

## erlang

The following functions are exported:

- `abs(Number)`  
[page 63]
- `apply({Module, Function}, ArgumentList)`  
[page 63]
- `apply(Module, Function, ArgumentList)`  
[page 63]
- `atom_to_list(Atom)`  
[page 63]
- `binary_to_list(Binary)`  
[page 63]
- `binary_to_list(Binary, Start, Stop)`  
[page 63]
- `binary_to_term(Binary)`  
[page 64]
- `erlang:cancel_timer(Ref)`  
[page 64]
- `erlang:check_process_code(Pid, Module)`  
[page 64]
- `concat_binary(ListOfBinaries)`  
[page 64]
- `date()`  
[page 64]
- `erlang:delete_module(Module)`  
[page 64]
- `erlang:demonitor(Ref)`  
[page 65]
- `erlang:disconnect_node(Node)`  
[page 65]
- `element(N, Tuple)`  
[page 65]
- `erase()`  
[page 65]
- `erase(Key)`  
[page 65]
- `exit(Reason)`  
[page 66]
- `exit(Pid, Reason)`  
[page 66]
- `float(Number)`  
[page 66]
- `float_to_list(Float)`  
[page 67]

- `get()`  
[page 67]
- `get(Key)`  
[page 67]
- `erlang:get_cookie()`  
[page 67]
- `get_keys(Value)`  
[page 67]
- `group_leader()`  
[page 68]
- `group_leader(Leader, Pid)`  
[page 68]
- `halt()`  
[page 68]
- `halt(Status)`  
[page 68]
- `erlang:hash(Term, Range)`  
[page 68]
- `hd(List)`  
[page 68]
- `erlang:info(What)`  
[page 68]
- `integer_to_list(Integer)`  
[page 69]
- `is_alive()`  
[page 69]
- `is_process_alive(Pid)`  
[page 69]
- `length(List)`  
[page 69]
- `link(Pid)`  
[page 69]
- `list_to_atom(CharIntegerList)`  
[page 69]
- `list_to_binary(ListOfIntegers)`  
[page 69]
- `list_to_float(AsciiIntegerList)`  
[page 70]
- `list_to_integer(AsciiIntegerList)`  
[page 70]
- `list_to_pid(AsciiIntegerList)`  
[page 70]
- `list_to_tuple(List)`  
[page 70]
- `erlang:load_module(Module, Binary)`  
[page 70]

- `erlang:localtime()`  
[page 71]
- `erlang:localtime_to_universaltime(DateTime)`  
[page 71]
- `make_ref()`  
[page 71]
- `erlang:module_loaded(Module)`  
[page 71]
- `erlang:monitor(Type, Item)`  
[page 72]
- `monitor_node(Node, Flag)`  
[page 72]
- `node()`  
[page 72]
- `node(Arg)`  
[page 73]
- `nodes()`  
[page 73]
- `now()`  
[page 73]
- `open_port(PortName, PortSettings)`  
[page 73]
- `pid_to_list(Pid)`  
[page 75]
- `erlang:port_info(Port, Item)`  
[page 75]
- `erlang:ports()`  
[page 75]
- `erlang:pre_loaded()`  
[page 76]
- `erlang:process_display(Pid, Type)`  
[page 76]
- `process_flag(Flag, Option)`  
[page 76]
- `process_info(Pid)`  
[page 76]
- `process_info(Pid, Item)`  
[page 77]
- `processes()`  
[page 77]
- `erlang:purge_module(Module)`  
[page 78]
- `put(Key, Value)`  
[page 78]
- `register(Name, Pid)`  
[page 78]

- `registered()`  
[page 78]
- `round(Number)`  
[page 79]
- `self()`  
[page 79]
- `erlang:send_after(Time, Pid, Msg)`  
[page 79]
- `erlang:set_cookie(Node, Cookie)`  
[page 79]
- `setelement(Index, Tuple, Value)`  
[page 79]
- `size(Item)`  
[page 79]
- `spawn(Module, Function, ArgumentList)`  
[page 80]
- `spawn(Node, Module, Function, ArgumentList)`  
[page 80]
- `spawn_link(Module, Function, ArgumentList)`  
[page 80]
- `spawn_link(Node, Module, Function, ArgumentList)`  
[page 80]
- `spawn_opt(Module, Function, ArgumentList, Options)`  
[page 80]
- `split_binary(Binary, Pos)`  
[page 81]
- `erlang:start_timer(Time, Pid, Msg)`  
[page 81]
- `statistics(Type)`  
[page 81]
- `erlang:system_flag(Flag, Value)`  
[page 82]
- `term_to_binary(Term)`  
[page 82]
- `throw(Any)`  
[page 82]
- `time()`  
[page 83]
- `tl(List)`  
[page 83]
- `trace(Pid, How, Flaglist)`  
[page 83]
- `trunc(Number)`  
[page 84]
- `tuple_to_list(Tuple)`  
[page 85]

- `erlang:universaltime()`  
[page 85]
- `erlang:universaltime_to_localtime(DateTime)`  
[page 85]
- `unlink(Pid)`  
[page 85]
- `unregister(Name)`  
[page 85]
- `whereis(Name)`  
[page 86]

## error\_handler

The following functions are exported:

- `undefined_function(Module, Func, ArgList) -> term()`  
[page 87] Called when an undefined function is encountered
- `undefined_lambda(Module, Fun, ArgList) -> term()`  
[page 87] Called when an undefined lambda (fun) is encountered

## error\_logger

The following functions are exported:

- `start() -> {ok, Pid} | {error, What}`  
[page 89] Starts the error logger event manager.
- `start_link() -> {ok, Pid} | {error, What}`  
[page 89] Starts the error logger event manager.
- `error_report(Report) -> ok`  
[page 89] Sends a standard error report event to the error logger.
- `error_report(Type, Report) -> ok`  
[page 90] Sends a user defined error report type event.
- `info_report(Report) -> ok`  
[page 90] Sends an information report to the error logger.
- `info_report(Type, Report) -> ok`  
[page 90] Sends a user defined information report type event.
- `error_msg(Format) -> ok`  
[page 91] Sends an error event to the error logger.
- `error_msg(Format, Args) -> ok`  
[page 91] Sends an error event to the error logger.
- `format(Format, Args) -> ok`  
[page 91] Sends an error event to the error logger.
- `info_msg(Format) -> ok`  
[page 91] Sends an information event to the error logger.

- `info_msg(Format, Args) -> ok`  
[page 91] Sends an information event to the error logger.
- `tty(Flag) -> ok`  
[page 91] Enables or disables error printouts to the tty.
- `logfile(Request) -> ok | FileName | {error, What}`  
[page 91] Enables or disables error printouts to a file.
- `add_report_handler(Module) -> ok | Other`  
[page 92] Adds a new event handler to the error logger.
- `add_report_handler(Module, Args) -> ok | Other`  
[page 92] Adds a new event handler to the error logger.
- `delete_report_handler(Module) -> Return | {error, What}`  
[page 92] Deletes an error report handler.
- `swap_handler(ToHandler) -> ok`  
[page 92] Swap from a primitive first handler to a standard event handler

## file

The following functions are exported:

- `read_file(Filename)`  
[page 94] Read a file
- `write_file(Filename, Binary)`  
[page 94] Write a file
- `get_cwd()`  
[page 95] Get the current working directory
- `get_cwd(Drive)`  
[page 95] Get the current working directory for the drive specified
- `set_cwd(DirName)`  
[page 95] Set the current working directory
- `delete(Filename)`  
[page 95] Delete a file
- `rename(Source, Destination)`  
[page 96] Rename a file
- `make_dir(DirName)`  
[page 96] Make a directory
- `del_dir(DirName)`  
[page 96] Delete a directory
- `make_link(Existing, New)`  
[page 97] Make a hard link to a file
- `make_symlink(Name1, Name2)`  
[page 97] Make a symbolic link to a file or directory
- `list_dir(DirName)`  
[page 97] List files in a directory
- `read_file_info(Filename)`  
[page 97] Get information about a file

- `read_link_info(Filename)`  
[page 98] Get information about a link or file
- `read_link(Linkname)`  
[page 99] See what a link is pointing to
- `file_info(Filename)`  
[page 99] Get information about a file
- `write_file_info(Filename, FileInfo)`  
[page 99] Change file information
- `change_owner(Filename, Uid)`  
[page 100] Change owner of a file
- `change_owner(Filename, Uid, Gid)`  
[page 100] Change owner for a file
- `change_group(Filename, Gid)`  
[page 100] Change owner for a file
- `change_time(Filename, Mtime)`  
[page 100] Change the modification time for a file
- `change_time(Filename, Mtime, Atime)`  
[page 100] Change the modification time for a file
- `consult(Filename)`  
[page 101] Read Erlang terms from a file
- `path_consult(Path, Filename)`  
[page 101] Read Erlang terms from a file
- `eval(Filename)`  
[page 101] Evaluate expressions in a file
- `path_eval(Path, Filename)`  
[page 101] Evaluate expressions in a file
- `open(Filename, ModeList)`  
[page 102] Open a file
- `path_open(Path, Filename, Mode)`  
[page 103] Open a file for access
- `close(IoDevice)`  
[page 103] Close a file
- `read(IoDevice, Number)`  
[page 103] Read from a file
- `write(IoDevice, Bytes)`  
[page 103] Write to a file
- `position(IoDevice, Location)`  
[page 104] Set position in a file
- `pread(IoDevice, Location, Number)`  
[page 104] Write to a file at a certain position
- `pwrite(IoDevice, Location, Bytes)`  
[page 104] Write to a file at a certain position
- `truncate(IoDevice)`  
[page 104] Truncate a file
- `sync(IoDevice)`  
[page 104] Synchronizes the in-memory state of a file with that on the physical medium
- `format_error(ErrorDescriptor)`  
[page 104] Returns an English description of an error term

## gen\_tcp

The following functions are exported:

- `accept(ListenSocket) -> {ok, Socket} | {error, Reason}`  
[page 108] Accepts an incoming connection request on a listen socket.
- `accept(ListenSocket, Timeout) -> {ok, Socket} | {error, Reason}`  
[page 108] Accepts an incoming connection request on a listen socket.
- `close(Socket) -> ok | {error, Reason}`  
[page 108] Closes an TCP socket
- `connect(Address, Port, Options) -> {ok, Socket} | {error, Reason}`  
[page 108] Connects to a TCP port.
- `connect(Address, Port, Options, Timeout) -> {ok, Socket} | {error, Reason}`  
[page 108] Connects to a TCP port.
- `controlling_process(Socket, NewOwner) -> ok | {error, eperm}`  
[page 109] Assigns a new controlling process to a socket
- `listen(Port, Options) -> {ok, Socket} | {error, Reason}`  
[page 109] Sets up a socket which listen on Port
- `recv(Socket, Length) -> {ok, Packet} | {error, Reason}`  
[page 109] Receives a packet from a passive socket
- `recv(Socket, Length, Timeout)`  
[page 109] Receives a packet from a passive socket
- `send(Socket, Packet) -> ok | {error, Reason}`  
[page 110] Sends a packet

## gen\_udp

The following functions are exported:

- `close(Socket) -> ok | {error, Reason}`  
[page 111] Close Socket.
- `controlling_process(Socket, NewOwner) ->`  
[page 111] Change controlling process of a Socket.
- `open(Port) -> {ok, Socket } | { error, Reason }`  
[page 111] Associates a UDP port number with the process calling it.
- `open(Port, Options) -> {ok, Socket } | { error, Reason }`  
[page 111] Associates a UDP port number with the process calling it.
- `recv(Socket, Length) -> {ok, Packet} | {error, Reason}`  
[page 112] Receives a packet from a passive socket
- `recv(Socket, Length, Timeout)`  
[page 112] Receives a packet from a passive socket
- `send(S, Address, Port, Packet) -> ok | {error, Reason}`  
[page 112] Sends a packet to a specified Address and Port (from port associated with Id).

## global

The following functions are exported:

- `del_lock(Id)`  
[page 113] Deletes the lock Id
- `del_lock(Id, Nodes) -> void()`  
[page 113] Deletes the lock Id
- `notify_all_name(Name, Pid1, Pid2) -> none`  
[page 114] Name resolving function that notifies both Pids
- `random_exit_name(Name, Pid1, Pid2) -> Pid1 | Pid2`  
[page 114] Name resolving function that kills one Pid
- `random_notify_name(Name, Pid1, Pid2) -> Pid1 | Pid2`  
[page 114] Name resolving function that notifies one Pid
- `register_name(Name, Pid)`  
[page 114] Globally registers Pid as Name
- `register_name(Name, Pid, Resolve) -> yes | no`  
[page 114] Globally registers Pid as Name
- `registered_names() -> [Name]`  
[page 114] Returns all globally registered names
- `re_register_name(Name, Pid)`  
[page 115] Atomically re-registers Pid for Name
- `re_register_name(Name, Pid, Resolve) -> void()`  
[page 115] Atomically re-registers Pid for Name
- `send(Name, Msg) -> Pid`  
[page 115] Sends Msg to the global process Name
- `set_lock(Id)`  
[page 115] Sets a lock on the specified nodes
- `set_lock(Id, Nodes)`  
[page 115] Sets a lock on the specified nodes
- `set_lock(Id, Nodes, Retries) -> boolean()`  
[page 115] Sets a lock on the specified nodes
- `start()`  
[page 116] Starts the global name server
- `start_link() -> {ok, Pid} | {error, Reason}`  
[page 116] Starts the global name server
- `stop() -> void()`  
[page 116] Stops the global name server
- `sync() -> void()`  
[page 116] Synchronizes the global name server
- `trans(Id, Fun)`  
[page 116] Micro transaction facility
- `trans(Id, Fun, Nodes)`  
[page 116] Micro transaction facility
- `trans(Id, Fun, Nodes, Retries) -> Res | aborted`  
[page 116] Micro transaction facility

- `unregister_name(Name) -> void()`  
[page 116] Unregisters the global name `Name`
- `whereis_name(Name) -> Pid() | undefined`  
[page 116] Returns the `Pid` of the global process `Name`

## global\_group

The following functions are exported:

- `global_groups() -> {OwnGroupName, [OtherGroupName]} | undefined`  
[page 118] Returns the global group names
- `info() -> [{state, State}, {own_group_name, atom()}, {own_group_nodes, [Node]}, {synced_nodes, [Node]}, {sync_error, [Node]}, {no_contact, [Node]}, {other_groups, Other_grps}, {monitoring, [pid()]}]`  
[page 118] Returns the state of the global group process
- `monitor_nodes(Flag) -> ok`  
[page 118] Subscription of node status for nodes in the immediate global group
- `own_nodes() -> [Node] | {error, ErrorMessage}`  
[page 118] Returns the global group names
- `registered_names({node, Node}) -> [Name] | {error, ErrorMessage}`  
[page 118] Returns all globally registered names
- `registered_names({group, GlobalGroupName}) -> [Name]`  
[page 118] Returns all globally registered names
- `send(Name, Msg) -> Pid | {badarg, Msg} | {error, ErrorMessage}`  
[page 119] Sends `Msg` to a registered process `Name`
- `send({node, Node}, Name, Msg) -> Pid | {badarg, Msg} | {error, ErrorMessage}`  
[page 119] Sends `Msg` to a registered process `Name`
- `send({group, GlobalGroupName}, Name, Msg) -> Pid | {badarg, Msg} | {error, ErrorMessage}`  
[page 119] Sends `Msg` to a registered process `Name`
- `sync() -> ok`  
[page 119] Synchronizes the immediate global group
- `whereis_name(Name) -> Pid | undefined | {error, ErrorMessage}`  
[page 119] Returns the `Pid` of the global process `Name`
- `whereis_name({node, Node}, Name) -> Pid | undefined | {error, ErrorMessage}`  
[page 119] Returns the `Pid` of the global process `Name`
- `whereis_name({group, GlobalGroupName}, Name) -> Pid | undefined | {error, ErrorMessage}`  
[page 119] Returns the `Pid` of the global process `Name`
- `start()`  
[page 120] Starts the global group server
- `start_link() -> {ok, Pid} | {error, Reason}`  
[page 120] Starts the global group server
- `stop() -> void()`  
[page 120] Stops the global group server

## heart

The following functions are exported:

- `start()` -> {ok, Pid} | ignore | {error, What}  
[page 121] Starts the heart program.
- `set_cmd(Cmd)` -> ok | {error, {bad\_cmd, Cmd}}  
[page 122] Sets a temporary reboot command.
- `clear_cmd()` -> ok  
[page 122] Clears the temporary boot command.

## inet

The following functions are exported:

- `format_error(Tag)`  
[page 123] Returns a diagnostic error string.
- `gethostbyaddr(Address)` -> {ok, Hostent} | {error, Reason}  
[page 123] Returns a hostent record for the host with the given address
- `gethostbyname(Name)` -> {ok, Hostent} | {error, Reason}  
[page 123] Returns a hostent record for the host with the given name
- `gethostbyname(Name, Family)` -> {ok, Hostent} | {error, Reason}  
[page 124] Returns a hostent record for the host with the given name
- `gethostname()` -> {ok, Name} | {error, Reason}  
[page 124] Returns the local hostname
- `sockname(Socket)` -> {ok, {IP, Port}} | {error, Reason}  
[page 124] Returns the local address and port number for a socket.
- `peername(Socket)` -> {ok, {Address, Port}} | {error, Reason}  
[page 124] Returns the address and port for the other end of a connection.
- `port(Socket)` -> {ok, Number}  
[page 124] Returns the local port number for a socket.
- `close(Socket)` -> ok  
[page 124] Closes a socket of any type
- `getaddr(IP, inet)` -> {ok, {A1, A2, A3, A4}} | {error, Reason}  
[page 125] Returns the IP-address for IP
- `setopts(Socket, Options)` -> ok | {error, Reason}  
[page 125] Sets one or more options for a socket.

## init

The following functions are exported:

- `boot(BootArgs)` -> void()  
[page 130] Start the Erlang system.

- `get_arguments()` -> Flags  
[page 130] Get all flag arguments.
- `get_argument(Flag)` -> {ok, Values} | error  
[page 131] Get values associated with an argument.
- `get_args()` -> [Arg]  
[page 131] Get all (non-flag) arguments.
- `restart()` -> void()  
[page 131]
- `reboot()` -> void()  
[page 131]
- `stop()` -> void()  
[page 131]
- `get_status()` -> {InternalStatus, ProvidedStatus}  
[page 131] Get status information during system start.
- `script_id()` -> Id  
[page 132] Get the identity of the used boot script.

## net\_adm

The following functions are exported:

- `host_file()`  
[page 134]
- `dns_hostname(Host)`  
[page 134]
- `localhost()`  
[page 134]
- `names(), names(Host)`  
[page 134]
- `ping(Node)`  
[page 134]
- `world (), world (verbose)`  
[page 134]
- `world_list (Hostlist), world_list (Hostlist, verbose)`  
[page 134]

## net\_kernel

The following functions are exported:

- `kernel_apply(M, F, A)`  
[page 136]
- `monitor_nodes(Flag)`  
[page 136]
- `allow(NodeList)`  
[page 136]

## OS

The following functions are exported:

- `cmd(Command) -> string()`  
[page 137] Executes Command in a command shell of the target OS.
- `find_executable(Name) -> Filename | false`  
[page 137] Returns the absolute filename of a program.
- `find_executable(Name, Path) -> Filename | false`  
[page 137] Returns the absolute filename of a program.
- `getenv(VarName) -> Value | false`  
[page 137] Returns the Value of the environment variable VarName.
- `getpid() -> Value`  
[page 138] Returns the process identifier of the emulator process as a string.
- `type() -> {Osfamily,Osname} | Osfamily`  
[page 138] Returns the Osfamily and, in some cases, Osname of the current operating system.
- `version() -> {Major, Minor, Release} | VersionString`  
[page 138] Returns the Operating System version.

## pg2

The following functions are exported:

- `create(Name) -> void()`  
[page 139] Creates a new, empty process group
- `delete(Name) -> void()`  
[page 139] Deletes a process group
- `get_closest_pid(Name) -> Pid | {error, Reason}`  
[page 139] Common dispatch function
- `get_members(Name) -> [Pid] | {error, Reason}`  
[page 139] Returns all processes in a group
- `get_local_members(Name) -> [Pid] | {error, Reason}`  
[page 140] Returns all local processes in a group
- `join(Name, Pid) -> ok | {error, Reason}`  
[page 140] Joins a process to a group
- `leave(Name, Pid) -> ok | {error, Reason}`  
[page 140] Makes a process leave a group
- `which_groups() -> [Name]`  
[page 140] Returns a list of all known groups
- `start()`  
[page 140] Starts the pg2 server
- `start_link() -> {ok, Pid} | {error, Reason}`  
[page 140] Starts the pg2 server

## rpc

The following functions are exported:

- `start()`  
[page 141]
- `stop()`  
[page 141]
- `call(Node, Module, Function, Args)`  
[page 141]
- `cast(Node, Module, Function, Args)`  
[page 141]
- `block_call(Node, Mod, Fun, Args)`  
[page 141]
- `server_call(Node, Name, ReplyWrapper, Msg)`  
[page 142]
- `abcast(Name, Mess)`  
[page 142]
- `abcast(Nodes, Name, Mess)`  
[page 142]
- `sbcast(Name, Msg)`  
[page 142]
- `sbcast(Nodes, Name, Msg)`  
[page 142]
- `eval_everywhere(Mod, Fun, Args)`  
[page 142]
- `eval_everywhere(Nodes, Mod, Fun, Args)`  
[page 142]
- `multicall(M, F, A)`  
[page 142]
- `multicall(Nodes, M, F, A)`  
[page 143]
- `multi_server_call(Name, Msg)`  
[page 143]
- `multi_server_call(Nodes, Name, Msg)`  
[page 143]
- `safe_multi_server_call(Name, Msg)`  
[page 143]
- `safe_multi_server_call(Nodes, Name, Msg)`  
[page 143]
- `async_call(Node, Mod, Fun, Args)`  
[page 144]
- `yield(Key)`  
[page 144]
- `nb_yield(Key)`  
[page 144]

- `parallel_eval(ListOfTuples)`  
[page 144]
- `pmap({M, F}, Extraargs, List)`  
[page 144]
- `pinfo(Pid)`  
[page 144]
- `pinfo(Pid, Item)`  
[page 144]

## seq\_trace

The following functions are exported:

- `set_token(Component, ComponentValue) -> {Component, PreviousValue}`  
[page 145] Sets the individual Component of the trace token.
- `set_token(Token) -> PreviousToken`  
[page 146] Sets the trace token to Value.
- `get_token(Component) -> {Component, ComponentValue}`  
[page 146] Returns the ComponentValue of the trace token component Component.
- `get_token() -> TraceToken`  
[page 146] Returns the value of the trace token.
- `print(TraceInfo) -> void`  
[page 146] Puts the Erlang term TraceInfo into the sequential trace output.
- `reset_trace() -> void`  
[page 147] Stops all sequential tracing on the Erlang node.
- `set_system_tracer(Pid) -> PreviousPid`  
[page 147] Sets the process Pid to be the system tracer.
- `get_system_tracer() -> pid() | false`  
[page 147] Returns the pid() of the current system tracer.

## user

The following functions are exported:

- `start() -> void()`  
[page 153] Starts the standard I/O system.

## wrap\_log\_reader

The following functions are exported:

- `chunk(Continuation)`  
[page 154] Reads a chunk of objects written to a wrap log.

- `chunk(Continuation, N) -> {Continuation2, Terms} | {Continuation2, Terms, Badbytes} | {Continuation2, eof} | {error, Reason}`  
[page 154] Reads a chunk of objects written to a wrap log.
- `close(Continuation) -> ok`  
[page 155] Closes a log
- `open(Filename) -> OpenRet`  
[page 155] Opens a log file
- `open(Filename, N) -> OpenRet`  
[page 155] Opens a log file

## **app**

No functions are exported

## **config**

No functions are exported

# kernel (Application)

The `kernel` application is the first application started, and it is one of two mandatory applications. The Kernel application contains the following services:

- `application_controller`
- `auth`
- `code`
- `disk_log`
- `erl_boot_server`
- `erl_ddll`
- `error_logger`
- `file`
- `global_group`
- `global_name_server`
- `net_kernel`
- `os`
- `rpc`
- `pg2`
- `timer`
- `user`

It is possible to synchronize a set of Erlang nodes. One can specify for a node to wait a specified amount of time for other nodes to become alive.

## Error Logger Event Handlers

Two error logger event handlers are defined in the kernel application. These are described in `error_logger(3)`.

## Configuration

The following configuration parameters are defined for the kernel application. See [application\(3\)](#) for more information about configuration parameters.

`distributed = [Distrib] <optional>` Specifies which applications are distributed and on which nodes they may execute. In this parameter:

- `Distrib = {AppName, Nodes} | {AppName, Time, Nodes}`
- `AppName = atom()`
- `Time = integer() > 0`
- `Nodes = [node() | {node(), ..., node()}]`

These parameters are described in [application\(3\)](#).

`permissions = [Perm] <optional>` Specifies the default permission for applications when they are started. In this parameter:

- `Perm = {AppName, Bool}`
- `AppName = atom()`
- `Bool = boolean()`

`error_logger = Value <optional>` Value is one of:

- `tty` All standard error reports are written to `stdio`. This is the default option.
- `{file, FileName}` All standard error reports are written to the file `FileName`, where `FileName` is a string.
- `false` No error logger handler is installed.

`global_groups = [GroupName, [Node]] <optional>` Specifies the groups of nodes which will have their own global name space. In this parameter:

- `GroupName = atom()`
- `Node = atom()`

These parameters are described in [global\\_group\(3\)](#).

`inet_parse_error_log = LogMode <optional>` `LogMode` is one of:

- `silent` No `error_logger` messages are generated when erroneous lines are found and skipped in the various configuration files. The default if the variable is not set is that erroneous lines are reported via the `error_logger`.

`net_ticktime = TickTime <optional>` Specifies the `net_kernel` tick time.

`TickTime` is given in seconds. Once every `TickTime / 4` second, all connected nodes are ticked (if anything else has been written to a node) and if nothing has been received from another node within the last four (4) tick times that node is considered to be down. This ensures that nodes which are not responding, for reasons such as hardware errors, are considered to be down.

The time `T`, in which a node that is not responding is detected, is calculated as:  $\text{MinT} < T < \text{MaxT}$  where

$$\begin{aligned} \text{MinT} &= \text{TickTime} - \text{TickTime} / 4 \\ \text{MaxT} &= \text{TickTime} + \text{TickTime} / 4 \end{aligned}$$

`TickTime` is by default 60 (seconds). Thus,  $45 < T < 75$  seconds.

*Note:* All communicating nodes should have the same `TickTime` value specified.

*Note:* Normally, a terminating node is detected immediately.

- `sync_nodes_mandatory` = [NodeName] <optional> Specifies which other nodes must be alive in order for this node to start properly. If some node in this list does not start within the specified time, this node will not start either. If this parameter is undefined, it defaults to the empty list.
- `sync_nodes_optional` = [NodeName] <optional> Specifies which other nodes can be alive in order for this node to start properly. If some node in this list does not start within the specified time, this node starts anyway. If this parameter is undefined, it defaults to the empty list.
- `sync_nodes_timeout` = integer() | infinity <optional> Specifies the amount of time (in milliseconds) this node will wait for the mandatory and optional nodes to start. If this parameter is undefined, no node synchronization is performed. This option also makes sure that global is synchronized.
- `start_ddll` = true | false <optional> Starts the `ddll_server` if the parameter is true (see `erl_ddll(3)`). This parameter should be set to true in an embedded system which uses this service.  
The default value is false.
- `start_dist_ac` = true | false <optional> Starts the `dist_ac` server if the parameter is true (see `application(3)`). This parameter should be set to true for systems that use distributed applications.  
The default value is false. If this parameter is undefined, the server is started if the parameter `distributed` is set.
- `start_boot_server` = true | false <optional> Starts the `boot_server` if the parameter is true (see `erl_boot_server(3)`). This parameter should be set to true in an embedded system which uses this service.  
The default value is false.
- `boot_server_slaves` = [SlaveIP] <optional> If the `start_boot_server` configuration parameter is true, this parameter can be used to initialize `boot_server` with a list of slave IP addresses. `SlaveIP` = string() | atom | { integer(), integer(), integer(), integer() }  
where  $0 \leq \text{integer()} \leq 255$ .  
Examples of `SlaveIP` in atom, string and tuple form are:  
'150.236.16.70', "150,236,16,70", {150,236,16,70}.  
The default value is [].
- `start_disk_log` = true | false <optional> Starts the `disk_log_server` if the parameter is true (see `disk_log(3)`). This parameter should be set to true in an embedded system which uses this service.  
The default value is false.
- `start_os` = true | false <optional> Starts the `os_server` if the parameter is true (see `os(3)`). This parameter should be set to true in an embedded system which uses this service.  
The default value is false.
- `start_pg2` = true | false <optional> Starts the `pg2` server (see `pg2(3)`) if the parameter is true. This parameter should be set to true in an embedded system which uses this service.  
The default value is false.
- `start_timer` = true | false <optional> Starts the `timer_server` if the parameter is true (see `timer(3)`). This parameter should be set to true in an embedded system which uses this service.  
The default value is false.

`keep_zombies = integer() <optional>` Sets the value of the system flag `keep_zombies`.  
The default value is 0.

## See Also

`application(3)`, `auth(3)`, `code(3)`, `disk_log(3)`, `erl_ddll(3)`, `erl_boot_server(3)`, `error_logger(3)`, `file(3)`, `global(3)`, `global_group(3)`, `net_kernel(3)`, `os(3)`, `pg2(3)`, `rpc(3)`, `timer(3)`, `user(3)`

# application (Module)

This module contains functions for controlling applications (eg. starting and stopping applications), and functions to access information about any application, (eg. configuration parameters)

All applications are started by the `application_controller` process. Each application has an `application_master` process. This process monitors the application and reports to the application controller if the application terminates.

An application can be started locally or distributed. A distributed application is started on one of several nodes while a local application is always started on the current node.

The local applications are controlled by the application controller. The distributed applications are controlled by another process, called the distributed application controller (`dist_ac`). The distributed application controller on different nodes monitor each other. Therefore, if a node goes down, the distributed applications on that node will be automatically re-started on one of the remaining nodes.

The distributed application controller is not started by default. Systems that use distributed applications must set the configuration parameter `start_dist_ac` in `kernel`.

## Exports

```
get_all_env()
```

```
get_all_env(Application) -> Env
```

Types:

- Application = atom()
- Env = [{Key, Value}]
- Key = atom()
- Value = term()

Retrieves the values of the application's configuration parameters. If `Application` is not specified, then the configuration parameters for the application which executes the call are returned.

```
get_all_key()
```

```
get_all_key(Application) -> {ok, Keys} | undefined
```

Types:

- Application = atom()
- Keys = [{Key, Value}]
- Key = atom()
- Value = term()

Retrieves all the keys from the application's resource file, `Application.app`. If `Application` is not specified, then the keys for the application which executes the call are returned.

```
get_application()
```

```
get_application(Pid | Module) -> {ok, Application} | undefined
```

Types:

- `Pid = pid()`
- `Module = atom()`
- `Application = atom()`

Retrieves the name of the application where the process `Pid` executes. If `Pid` is not specified, `self()` is used. If an atom is given the name of the application which contains the module will be returned, or `undefined`.

```
get_env(Key)
```

```
get_env(Application, Key) -> {ok, Value} | undefined
```

Types:

- `Application = atom()`
- `Key = atom()`
- `Value = term()`

Retrieves the value of an application's configuration parameter. If `Application` is not specified, the parameter for the application which executes the call is retrieved.

```
get_key(Key)
```

```
get_key(Application, Key) -> {ok, Value} | undefined
```

Types:

- `Application = atom()`
- `Key = atom()`
- `Value = term()`

Retrieves the key from the application's resource file, `Application.app`. If `Application` is not specified, then the key for the application which executes the call is returned.

```
load(Application)
```

```
load(Application, DistNodes) -> ok | {error, Reason}
```

Types:

- `Application = atom() | appl_descr()`
- `DistNodes = {Name, Nodes} | {Name, Time, Nodes} | default`
- `appl_descr() = {application, Name, [appl_opt()]}`
- `Name = atom()`
- `Time = integer() > 0`
- `Nodes = [node() | {node(), ..., node()}]`
- `appl_opt() = {description, string()} | {vsn, vsn()} | {modules, [{atom(), vsn()}]} | {registered, [atom()}] | {applications, [atom()}] | {env, [{atom(), term()}]} | {mod, {Mod, StartArgs}}`

- `vsn() = string()`

If the name of the application is given, the application controller searches the current path (the same as the code path) for a file called `Application.app`.

*Note:* This file must contain the `appl_descr()` (written in plain text, with a dot and space after the term).

`description` and `version` - Contains information about an application that can be retrieved by calling `application:loaded_applications/0`.

`modules` - Lists the modules that this application introduces.

`registered` is a list of the registered names that this application uses for its own processes.

`applications` - Lists of other applications that must be started before this one.

`env` is a list of configuration parameters. *Note:* The definitions in this list may be altered by definitions in the system configuration file, specified by the command line argument `-config`. They can also be altered directly from the command line, by giving `-Name Par Value`.

`mod` is the application call back module. `Mod:start(StartType, StartArgs)` is called when the application is started. Refer to the call back function `start/2`.

The `DistNodes` parameter will override the value of the application in the kernel configuration parameter `distributed`. The data structure specifies a list of nodes where the application `Name` may execute. If the nodes are specified in a tuple, the order of where to start the application will be undefined. If a node crashes and `Time` has been specified, then the application controller will wait for `Time` milliseconds before attempting to restart the application on another node. If `Time` is not specified, it will default to 0. If a node goes down, the application will be restarted immediately on another node. If `DistNodes` is default, the value in the configuration parameter `distributed` will be used.

```
loaded_applications() -> [{Name, Description, Version}]
```

Types:

- `Name = atom()`
- `Description = string()`
- `Version = string()`

This function returns a list of applications which are loaded in the system. `Description` and `Version` are as defined in the application specification.

```
permit(Application, Bool) -> ok | {error, Reason}
```

Types:

- `Name = atom()`
- `Bool = bool()`

This function changes an application's permission to run on the node, or vice versa. If the permission of a locally running application is set to `false`, the application will be stopped. When the permission is set to `true`, the local application will be started. If the permission of a running, distributed application is set to `false`, the application will be moved to another node where it may run, if a node is available.

The application must be loaded before the permit function can be called.

This function does not return until the application is either started, stopped or successfully moved to another node. However, in some cases where permission is set to `true` the function may return `ok` even though the application itself has not started. This is true when an application cannot start because it has dependencies on applications which have not yet been started. When these applications are started the dependent application will also be started.

By default, all applications are loaded with permission `true` on all nodes. The permission is configurable with the parameter `permissions` in `kernel`.

```
start(Application)
```

```
start(Application, Type) -> ok | {error, Reason}
```

Types:

- Application = `atom()`
- Type = `permanent` | `transient` | `temporary`

This function starts and application. If the application is not loaded, the application controller will first try to load it, as if `application:load(Application)` was called.

The Type specifies what happens if the application dies.

- If a permanent application dies, all other applications are also terminated.
- If a transient application dies normally, this is reported and no other applications are terminated. If a transient application dies abnormally, all other applications are also terminated.
- If a temporary application dies this is reported and no other applications are terminated. In this way, an application can run in test mode, without disturbing the other applications.

Default value for Type is `temporary`.

```
start_type() -> normal | local | {takeover, node()} | {failover, node()}
```

This function returns the type of application start which is executing.

`normal` is returned when an application is starting and the below circumstances have not occurred.

`local` is returned if a supervised process restarts due to abnormal exit or if no start is running at the time of request.

`{takeover, Node}` is returned if the application is requested to move to another node either due to a call to `takeover/2` or when a node with higher priority to run the application is restarted.

`{failover, Node}` is returned if the application is restarted due to the Node crashing where the application was previously executing.

```
stop(Application) -> ok
```

Types:

- Application = atom()

This function stops a running application. If the application was distributed, no other node will restart it.

```
takeover(Application, Type) -> {ok, Pid} | {error, Reason}
```

Types:

- Application = atom()
- Type = permanent | transient | temporary

This function moves a distributed application which executes on another node *Node* to the current node. The application is started by calling `Mod:start({takeover, Node}, StartArgs)` before the application is stopped on the other node. This makes it possible to transfer application specific data from a currently running application to a new node. When the application start function returns, the application on a *Node* is stopped. This means that two instances of the application may be running on two different nodes at one time. If this is not acceptable, parts of the application on the old node (*Node*) may be shut down when the new node starts the application. *Note:* that the old application must not be stopped entirely (i.e. `application:stop/1` must not be called on the old node). The main supervisor, must still be alive.

```
which_applications() -> [{Name, Description, Version}]
```

Types:

- Name = atom()
- Description = string()
- Version = string()

Returns a list of the applications which are running in the system. *Description* and *Version* are as defined in the application specification.

## Call back Module

The following functions are exported from an `application` call back module.

## Exports

```
Module:config_change(Changed, New, Removed) -> ok
```

Types:

- Changed = [{Parameter, NewValue}]
- New = [{Parameter, Value}]
- Removed = [Parameter]
- Parameter = atom()

- `NewValue = term()`
- `Value = term()`

After an installation of a new release all started applications on a node are notified of the changed, new and removed configuration parameters. The unchanged configuration parameters are not affected and therefore the function is not evaluated for applications which have unchanged configuration parameters between the old and new releases.

```
Module:start(Type, ModuleStartArgs) -> {ok, Pid} | {ok, Pid, State} | {error, Reason}
```

Types:

- `Type = normal | {takeover, node()} | {failover, node()}`
- `ModuleStartArgs = term()`
- `Pid = pid()`
- `State = state()`

This function starts a primary application. Normally, this function starts the main supervisor of the primary application.

If `Type` is `{takeover, Node}`, it is a distributed application which is running on the `Node`. If the application does not have the `start-phases` key defined in the application's resource file, the application will be stopped by the application controller after this call returns (see `start-phase/3`) This makes it possible to transfer the internal state from the running application to the one to be started. This function must not stop the application on `Node`, but it may shut down parts of it. For example, instead of stopping the application, the main supervisor may terminate all its children.

If `Type` is `{failover, Node}`, the application will be restarted due to a crash of the node where the application was previously executing.

`{failover, node()}` is valid only if the `start_phases` key is defined in the application's resource file. Otherwise the type is set to `normal` at failover.

The `ModuleStartArgs` parameter is specified in the application resource file (`.app`), as `{mod, {Module, ModuleStartArgs}}`.

`State` is any term. It is passed to `Module:stop/1`. If no `State` is returned, `[]` is used.

```
Module:start_phase(Phase, Type, PhaseStartArgs) -> ok | {error, Reason}
```

Types:

- `Phase = atom()`
- `Type = normal | {takeover, node()} | {failover, node()}`
- `PhaseStartArgs = term()`
- `Pid = pid()`
- `State = state()`

This function starts a application in the phase `Phase`. It is called by default only for a primary application and not for the included applications, refer to User's Guide chapter 'Design Principles' regarding incorporating included applications.

The `PhaseStartArgs` parameter is specified in the application's resource file (`.app`), as `{start_phases, [{Phase, PhaseStartArgs}]}`, the `Module` as `{mod, {Module, ModuleStartArgs}}`.

This call back function is only valid for applications with a defined `start_phases` key. This function will be called once per `Phase`.

If `Type` is `{takeover, Node}`, it is a distributed application which runs on the `Node`. When this call returns for the last start phase, the application on `Node` will be stopped by the application controller. This makes it possible to transfer the internal state from the running application. When designing the start phase function it is imperative that the application is not allowed to terminate the application on `node`. However, it is possible to partially shut it down for eg. the main supervisor may terminate all the application's children.

If `Type` is `{failover, Node}`, due to a crash of the node where the application was previously executing, the application will restart.

```
Module:stop(State) -> void()
```

Types:

- `State = state()`

This function is called when the application has stopped, either because it crashed, or because someone called `application:stop`. It cleans up after the `Module:start/2` function.

`State` is the state that was returned from `Mod:start/2`, or `[]` if no state was returned.

## See Also

`kernel(3)`

## auth (Module)

Authentication determines which nodes are allowed to communicate with each other. In a network of different Erlang nodes, it is built into the system at the lowest possible level. Each node has its Magic Cookie, which is an Erlang atom.

Whenever a message is transferred from one node to another, it is accompanied by the Magic Cookie of the receiving node. For example, a message transferred from node A to node B is accompanied by what node A believes to be the Magic Cookie of node B.

When the message arrives at node B, the runtime system immediately checks that the accompanying cookie is the right one. If it is, the message is passed on in the normal way. If it is not, the message is transformed into a `badcookie` message, which is sent to the system process `net_kernel`. By default, the `net_kernel` process passes the message to the registered process `auth`, which is then responsible for taking the appropriate action for the unauthorized message. In the standard system, the default action is to shut down connection to that node.

At start-up, the first action of the standard `auth` server is to read a file named `$HOME/erlang.cookie`. An atom is created from the contents of this file and the cookie of the node is set to this atom with the use of `erlang:set_cookie(node(), CookieAtom)`.

If the file does not exist, it is created. The UNIX permissions mode of the file is set to octal 400 (read-only by owner) and filled with a random string. For this reason, the same user, or group of users with identical cookie files, can have Erlang nodes which can communicate freely and without interference from the Magic Cookie system. Users who want to run nodes on separate file systems must be certain that their cookie files are identical on the different file systems.

Initially, each node has a random atom assigned as its magic cookie. Once the procedure described above has been concluded, the cookie is set to the contents of the `$HOME/erlang.cookie` file.

To communicate with another node, the magic cookie of that node must be known. The BIF `erlang:set_cookie(Node, Cookie)` sets the cookie for `Node` to `Cookie`. From then on, all messages will be accompanied by the cookie `Cookie`. If the cookie is not correct when messages arrive at `Node`, they are sent to the `auth` server at `Node`. The call `erlang:set_cookie(node(), CookieAtom)` will set the current cookie to `CookieAtom`. It will, however, also set the cookie of all other unknown nodes to `CookieAtom`. In the case of the default `auth` server, this is the first thing done when the system starts. The default then, is to assume that all nodes which communicate have the same cookie. In the case of a single user on a single file system, this is indeed true and no further action is required. The original cookie can also be fetched by the BIF `erlang:get_cookie()`.

If nodes which communicate do not have the same cookie, they can be set explicitly on each node with the aid of `erlang:set_cookie(Node, Cookie)`. All messages sent to the node `Node` will then be accompanied by the cookie `Cookie`. Distributed systems with multiple User IDs can be handled in this way.

Initially, the system cookie is set to a random atom, and the (assumed) cookie of all other nodes is initially set to the atom `nocookie`. Thus, the Erlang system is completely

unprotected when `erlang:set_cookie(node(), nocookie)` is run. Sometimes, this may be appropriate for systems which are not normally networked, and it can also be appropriate for maintenance purposes.

If a user wants to write a specific application with another authentication procedure than the one provided by the standard system, passwords for example, the auth server can be replaced by a user defined server. Once this user defined authentication procedure succeeds and the two nodes can exchange cookies, both nodes will use the `set_cookie/2` BIF to set the cookie of other nodes in the network so there can be communication between them.

In the standard system, the default when two nodes are connected is to immediately connect all other involved nodes as well. This way, there is always a fully connected network. If there are nodes with different cookies, this method might be inappropriate and the host OS command line option `-connect_all false` must be issued to the Erlang system. See `global(3)`.

This module uses the two BIFs `erlang:get_cookie()` which returns the magic cookie of the local node, and `erlang:set_cookie(Node, Cookie)` which sets the magic cookie of `Node` to `Cookie`. If `Node` is the user's node, the cookie of all other unknown nodes are also set to `Cookie` by this BIF.

## Exports

`start()`

Starts the auth server.

`stop()`

Stops the auth server.

`open(Name)`

This function opens up the server with the name `Name`. If, for example, node `N` is run with the cookie `C`, it is impossible for other nodes with other cookies to communicate with node `N`. The call `open/1` opens the server with the registered name `Name` so it can be accessed by any other node, irrespective of cookie. The call must be executed on both nodes to have any effect. All messages to the server must have the form `Name ! Msg` and all replies from the server `{Name, Reply}`, or `{Name, Node, Reply}`. With this feature, it is possible to perform specific tasks on publicly announced Erlang network servers.

`is_auth(Node)`

Returns the value `yes` if communication with `Node` is authorized, `no` if `Node` does not exist or communication is not authorized.

`exists(Node)`

Returns `yes` if `Node` exists, otherwise `no`.

`cookie()`

Reads `cookie` from `$HOME/.erlang.cookie` and sets it. This function is used by the `auth` server at start-up.

`node_cookie(Node, Cookie)`

If the cookie of `Node` is known to the user as `Cookie` but the user's cookie is not known at `Node`, this function informs `Node` of the identity of the user's cookie.

`node_cookie([Node, Cookie])`

Another version of the previous function with the arguments in a list which can be given on the host OS command line.

`cookie([Cookie])`

Equivalent to `erlang:set_cookie(node(), Cookie)`, but with the argument in a list so it can be given on the host OS command line.

# code (Module)

This module deals with the loading of compiled and interpreted code into a running Erlang system.

The code server dynamically loads modules into the system on demand, which means the first time the module is referenced. This functionality can be turned off using the command line flag `-mode embedded`. In this mode, all code is loaded during system start-up.

If started in *interactive* mode, all directories under the `$ROOT/lib` directory are initially added to the search path of the code server (`code:root_dir()`). The `$ROOT` directory is the installation directory of the Erlang system, `code:root_dir()`. Directories can be named `Name[-Vsn]` and the code server, by default, chooses the greatest (`>`) directory among those which have the same `Name`. The `-Vsn` suffix is optional.

If an `ebin` directory exists under a chosen directory, it is added to the directory. The `Name` of the directory (or library) can be used to find the full directory name (including the current version) through the `priv_dir/1` and `lib_dir/1` functions.

## Exports

```
start() -> {ok, Pid} | {error, What}
start(Flags) -> {ok, Pid} | {error, What}
```

Types:

- `Flags` = [`stick` | `nostick` | `embedded` | `interactive`]
- `Pid` = `pid()`
- `What` = `term()`

This function starts the code server. `start/0` implies that the `stick` and `interactive` flags are set.

`Flags` can also be entered as the command line flags `-stick`, `-nostick` and `-mode embedded | interactive`. `-stick` and `-mode interactive` are the defaults. The `stick` flag indicates that a module can never be re-loaded once it has been loaded from the kernel, `stdlib`, or `compiler` directories.

```
start_link() -> {ok, Pid} | {error, What}
start_link(Flags) -> {ok, Pid} | {error, What}
```

Types:

- `Flags` = [`stick` | `nostick` | `embedded` | `interactive`]
- `Pid` = `pid()`
- `What` = `term()`

This function starts the code server and sets up a link to the calling process. This function should be used if the code server is supervised. `start_link/0` implies that the `stick` and `interactive` flags are set.

The `Flags` can also be given as command line flags, `-stick`, `-nostick` and `-mode embedded | interactive` where `-stick` and `-mode interactive` is the default. The `stick` flag indicates that a module which has been loaded from the `kernel`, `stdlib` or compiler directories can never be reloaded.

```
set_path(DirList) -> true | {error, What}
```

Types:

- `DirList = [Dir]`
- `Dir = string()`
- `What = bad_directory | bad_path`

Sets the code server search path to the list of directories `DirList`.

```
get_path() -> Path
```

Types:

- `Path = [Dir]`
- `Dir = string()`

Returns the current path.

```
add_path(Dir) -> true | {error, What}
```

```
add_pathz(Dir) -> true | {error, What}
```

Types:

- `Dir = string()`
- `What = bad_directory`

Adds `Dir` to the current path. The directory is added as the last directory in the new path. If `Dir` already exists in the path, it is not added.

```
add_patha(Dir) -> true | {error, What}
```

Types:

- `Dir = string()`
- `What = bad_directory`

This function adds `Dir` to the beginning of the current path. If `Dir` already exists, the old directory is removed from path.

```
add_paths(DirList) -> ok
```

```
add_pathsz(DirList) -> ok
```

Types:

- `DirList = [Dir]`
- `Dir = string()`

This function adds the directories in `DirList` to the end of the current path. If a `Dir` already exists in the path, it is not added. This function always returns `ok`, regardless of the validity of each individual `Dir`.

```
add_pathsa(DirList) -> ok
```

Types:

- `DirList = [Dir]`
- `Dir = string()`

Adds the directories in `DirList` to the beginning of the current path. If a `Dir` already exists, the old directory is removed from the path. This function always returns `ok`, regardless of the validity of each individual `Dir`.

```
del_path(NameDir) -> true | false | {error, What}
```

Types:

- `NameDir = Name | Dir`
- `Name = atom()`
- `Dir = string()`
- `What = bad_name`

This function deletes an old occurrence of a directory in the current path with the name `.../Name[-*] [/ebin]`. It is also possible to give the complete directory name `Dir` in order to delete it.

This function returns `true` if the directory was deleted, and `false` if the directory was not found.

```
replace_path(Name, Dir) -> true | {error, What}
```

Types:

- `Name = atom()`
- `Dir = string()`
- `What = bad_name | bad_directory | {badarg, term()}`

This function replaces an old occurrence of a directory named `.../Name[-*] [/ebin]`, in the current path, with `Dir`. If `Name` does not exist, it adds the new directory `Dir` last in path. The new directory must also be named `.../Name[-*] [/ebin]`. This function should be used if a new version of the directory (library) is added to a running system.

```
load_file(Module) -> {module, Module} | {error, What}
```

Types:

- `Module = atom()`
- `What = nofile | sticky_directory | badarg | term()`

This function tries to load the Erlang module `Module`, using the current path. It looks for the object code file which has a suffix that corresponds to the Erlang machine used, for example `Module.jam`. The loading fails if the module name found in the object code differs from the name `Module`. `load_binary/3` must be used to load object code with a module name that is different from the file name.

```
load_abs(File) -> {module, Module} | {error, What}
```

Types:

- File = atom() | string()
- Module = atom()
- What = nofile | sticky\_directory | badarg | term()

This function does the same as `load_file(Module)`, but `File` is either an absolute file name, or a relative file name. The current path is not searched. It returns a value in the same way as `load_file(Module)`. Note that `File` should not contain an extension ("`.jam`" or "`.beam`"); `load_abs/1` adds the correct extension itself.

`ensure_loaded(Module) -> {module, Module} | {error, What} | {interpret, Module}`

Types:

- Module = atom()
- What = nofile | sticky\_directory | embedded | badarg | term()

This function tries to ensure that the module `Module` is loaded. To work correctly, a file with the same name as `Module.Suffix` must exist in the current search path. `Suffix` must correspond to the running Erlang machine, for example `.jam`. It returns a value in the same way as `load_file(File)`, or `{interpret, Module}` if `Module` is interpreted.

If the system is started with the `-mode embedded` command line flag, this function will not load a module which has not already been loaded. `{error, embedded}` is returned.

`delete(Module) -> true | false`

Types:

- Module = atom()

This function deletes the code in `Module` and the code in `Module` is marked as old. This means that no external function calls can be made to this occurrence of `Module`, but a process which executes code inside this module continues to do so. Returns `true` if the operation was successful (i.e., there was a current version of the module, but no old version), otherwise `false`.

`purge(Module) -> true | false`

Types:

- Module = atom()

This function purges the code in `Module`, that is, it removes code marked as old. If some processes still execute code in the old occurrence of `Module`, these processes are killed before the module is purged. Returns `true` if a process has been killed, otherwise `false`.

`soft_purge(Module) -> true | false`

Types:

- Module = atom()

This function purges the code in `Module`, that is, it removes code marked as old, but only if no process currently runs the old code. It returns `false` if a process uses the old code, otherwise `true`.

`is_loaded(Module) -> {file, Loaded} | false`

Types:

- Module = atom()
- Loaded = AbsFileName | preloaded | interpreted
- AbsFileName = string()

This function tests if module `Module` is loaded. If the module is loaded, the absolute file name of the file from which the code was obtained is returned.

`all_loaded()` -> [LoadMod]

Types:

- LoadMod = {Module, Loaded}
- Module = atom()
- Loaded = AbsFileName | preloaded | interpreted
- AbsFileName = string()

This function returns a list of tuples of the type {Module, Loaded} for all loaded modules. Loaded is the absolute file name of the loaded module, the atom `preloaded` if the module was pre-loaded, or the atom `interpreted` if the module is interpreted.

`load_binary(Module, File, Binary)` -> {module, Module} | {error, What}

Types:

- Module = atom()
- What = sticky\_directory | badarg | term()

This function can be used to load object code on remote Erlang nodes. It can also be used to load object code where the file name and module name differ. This, however, is a very unusual situation and should be used with care. The parameter `Binary` must contain object code for the module `Module`. The `File` parameter is only used by the code server to keep a record from which file the object code in `Module` comes. Accordingly, `File` is not opened and read by the code server.

`stop()` -> stopped

Stops the code server.

`root_dir()` -> RootDir

Types:

- RootDir = string()

Returns the root directory of the Erlang system, which is the directory where the Erlang system is installed.

`lib_dir()` -> LibDir

Types:

- LibDir = string()

Returns the library directory.

`lib_dir(Name)` -> LibDir | {error, What}

Types:

- Name = atom()

- LibDir = string()
- What = bad\_name

This function returns the current lib directory for the Name[-\*] directory (or library). The current path is searched for a directory named .../Name-\* (the -\* suffix is optional for directories in the search path and it represents the version of the directory).

compiler\_dir() -> CompDir

Types:

- CompDir = string()

This function returns the compiler directory.

priv\_dir(Name) -> PrivDir | {error, What}

Types:

- Name = atom()
- PrivDir = string()
- What = bad\_name

This function returns the current priv directory for the Name[-\*] directory. The current path is searched for a directory named .../Name-\* (the -\* suffix is optional for directories in the search path and it represents the version of the directory). The /priv suffix is added to the end of the found directory.

get\_object\_code(Module) -> {Module, Bin, AbsFileName} | error

Types:

- Module = atom()
- Bin = binary()
- AbsFileName = string()

This function searches the code path in the code server for the object code of the module Module. It returns {Mod, Bin, Filename} if successful, and error if not. Bin is a binary data object which contains the object code for the module. This can be useful if code is to be loaded on a remote node in a distributed system. For example, loading module Module on node N is done as follows:

```
...
{Mod, B, F} = code:get_object_code(Mod),
rpc:call(N,code,load_binary,[Mod,F,B]),
...
```

objfile\_extension() -> Ext

Types:

- Ext = string()

This function returns the object code file extension for the running Erlang machine, for example ".jam".

stick\_dir(Dir) -> ok | {error, term()}

Types:

- `Dir = string()`

This function marks `Dir` as 'sticky'. The system issues a warning and rejects the request if a user tries to re-load a module in a sticky directory. Sticky directories are used to warn the user about inadvertent changes to system software.

`unstick_dir(Dir) -> ok | {error, term()}`

Types:

- `Dir = string()`

This function unsticks a directory which has been marked sticky. Code which is located in the unstuck directory can be re-loaded into the system.

`which(Module) -> WhichFile`

Types:

- `Module = atom()`
- `WhichFile = FileName | non_existing | preloaded | interpreted`
- `FileName = string()`

If the module is not loaded already, this function returns the directory path to the first file name in the search path of the code server which contains the object code for `Module`. If the module is loaded, it returns the directory path to the file name which contains the loaded object code. If the module is pre-loaded or interpreted, this is returned instead. `non_existing` is returned if the module cannot be found.

`clash() -> ok`

Searches the entire code space for module names with identical names and writes a report to `stdout`.

`interpret(Module) -> {module, Module} | {error, What}`

Types:

- `Module = atom()`
- `What = no_interpreter | sticky_directory | badarg`

Marks `Module` as being interpreted.

`interpret_binary(Module, File, Binary) -> {module, Module} | {error, What}`

Types:

- `Module = atom()`
- `File = string()`
- `Binary = binary()`
- `What = no_interpreter | sticky_directory | badarg | term()`

Loads the interpreted `Module` into the interpreter. The parameter `Binary` contains the abstract form (and the source code) of the module. The file `File` parameter locates the used source code file.

`delete_interpret(Module) -> ok | {error, What}`

Types:

- Module = atom()
  - What = no\_interpreter | badarg
- Stops interpretation of Module.

interpreted() -> Modules

Types:

- Modules = [Module]
- Module = atom()

Returns a list of all modules which are being interpreted.

interpreted(Module) -> true | false

Types:

- Module = atom()

Returns true if Module is being interpreted, otherwise false.

## Notes

Dir has the described type `string()` in all functions. For backwards compatibility, `atom()` is also allowed, but `string()` is recommended.

The described type for `Module` is `atom()` in all functions. For backwards compatibility, `string()` is also allowed.

# disk\_log (Module)

`disk_log` is a disk based term logger. It is possible to efficiently log items to files. Each item which is logged is basically appended to a file. Two types of logs are supported, halt logs and wrap logs. For reasons of efficiency, items are always written to the file as binaries.

Two formats of the log files are supported, `internal` and `external`. The former uses an internally defined format for the log items. This format makes it possible to perform automatic repair of corrupt log files, and also makes it possible to read the logged terms from the file in a very efficient manner. The latter format leaves the format to the user of the log. If this format is used, the automatic repair and efficient reading of logs cannot be used.

A log file can be opened and closed. Whenever we open a non-existent log file, a new log file will be created. Logs using the `internal` format must be properly closed. If we try to open a log file which has not been properly closed, the `disk_log` module will automatically try to repair the log file by searching the file for magic bytes and Erlang terms.

When using the `internal` format for logs, the functions `log/2`, `log_terms/2`, `alog/2`, and `alog_terms/2` should be used. These functions take Erlang terms as argument for logging. When using the `external` format, the corresponding functions are `blog/2`, `blog_terms/2`, `balog/2`, and `balog_terms/2`. These functions take a list of bytes, or a binary with a list of bytes, as argument for logging. For example, to log the string "hello" in ASCII format, we can use `disk_log:blog(Log, "hello")`, or `disk_log:blog(Log, list_to_binary("hello"))`. These two alternatives are equally efficient. The `blog/2` functions can be used for `internal` formatted logs as well, but in this case they must be called with a binary constructed with a call to `term_to_binary/1`. There is no check to ensure this, it is entirely the responsibility of the caller. If these functions are called with a binary that does not correspond to an Erlang term, the `chunk/2,3` and automatic repair functions will fail. The corresponding term (not the binary) will be returned when `chunk/2,3` is called.

Logs can be configured to log to files on several nodes at the same time. In this case, the log is globally accessible. It is not guaranteed that all log files contain the same log items. This functionality ensures that as long as at least one of the involved nodes is alive at each time, all log items will be logged.

## Exports

```
alog(Log, Term) -> ok | {error, Reason}
balog(Log, Bytes) -> ok | {error, Reason}
```

Types:

- Log = term()
- Term = term()
- Bytes = binary() | [Byte]
- Byte = [Byte] | 0 =< integer() =< 255
- Reason = no\_such\_log

Asynchronously appends an item to the log. Accordingly, this function does not wait for the log process to actually write the object to the file. If the log is opened in read only mode the log attempt naturally fails, but the owner of the log will be notified of the failure only if the log is opened with `notify` set to `true`, see `open/1`.

The `alog` function is used for logs with format `internal`, and `balog` with format `external`.

```
alog_terms(Log, TermList) -> ok | {error, Reason}
```

```
balog_terms(Log, BytesList) -> ok | {error, Reason}
```

Types:

- Log = term()
- TermList = [term()]
- BytesList = [Bytes]
- Bytes = binary() | [Byte]
- Byte = [Byte] | 0 =< integer() =< 255
- Reason = no\_such\_log

Asynchronously appends a list of items to the log. If the log is opened in read only mode the log attempt naturally fails, but the owner of the log will be notified of the failure only if the log is opened with `notify` set to `true`, see `open/1`.

The `alog_terms` function is used for logs with format `internal`, and `balog_terms` with format `external`.

```
block(Log)
```

```
block(Log, QueueLogRecords) -> ok | {error, Reason}
```

Types:

- Log = term()
- QueueLogRecords = bool()

This function blocks a log. This means that all attempts to use the Log are suspended until the log is unblocked. The process that blocks a log may, however, use the function `chunk/2,3`. If `QueueLogRecords` is `true`, log attempts are suspended until the log is unblocked. If it is `false`, log records are discarded. Default is `true`.

```
change_size(Log, Size) -> ok | {error, Reason}
```

Types:

- Log = term()
- Size = integer() > 0 | infinity | {MaxBytes, MaxFiles}
- MaxBytes = integer() > 0
- MaxFiles = integer() > 0

This function changes the size of an open log. For a halt log it is always possible to increase the size, but it is not possible to decrease the size to be less than the current size of the file.

For a wrap log it is always possible to increase both the size and number of files. If the max files is decreased the change will not be valid until the current file is full and the log wraps to the next file. The redundant files will be removed next time the log wraps around, i.e. starts to log to the log number 1.

Example: the old max files was 10 and the new max files is 6. If the current file is not greater than the new max files the files 7-10 will be removed when the file 6 is full and the log starts to log to file number 1 again.

Otherwise the files greater than the current file will be removed when the current file is full (e.g. if the current file is 8, the files 9 and 10); the files between new max files and the current file (i.e. files 7 and 8) will be removed next time the new max log is full.

If the size of the files is decreased the change will immediately have affect on the current log. It won't of course change the size of logs already full until the next time they will be used.

If the log size is decreased for instance to save space, the function `inc_wrap_log/1` can be used to force the log to wrap.

`chunk(Log, Continuation)`

`chunk(Log, Continuation, N) -> {Continuation2, Terms} | {Continuation2, Terms, Badbytes} | eof | {error, Reason}`

Types:

- Log = `term()`
- Continuation = `start` | `cont()`
- N = `int() > 0` | `infinity`
- Continuation2 = `cont()`
- Terms= [`term()`]
- Badbytes = `integer()`

This function makes it possible to efficiently read the terms which have been appended to a log. It minimises disk I/O by reading large 8K chunks from the file.

The first time `chunk` is called an initial continuation, the atom `start`, must be provided.

When `chunk/3` is called, N controls the maximum number of terms that are read from the log in each chunk. Default is `infinity`, which means that all the terms contained in the 8K chunk are read. If less than N terms are returned, this does not necessarily mean that end of file is reached.

The `chunk` function returns a tuple `{Continuation2, Terms}`, where `Terms` is a list of terms found in the log. `Continuation2` is yet another continuation which must be passed on into any subsequent calls to `chunk`. With a series of calls to `chunk` it is then possible to extract all terms from a log.

The `chunk` function returns a tuple `{Continuation2, Terms, Badbytes}` if the log is opened in read only mode and the read chunk is corrupt. `Badbytes` indicates the number of non-Erlang terms found in the chunk. Note also that the log is not repaired.

`chunk` returns `eof` when the end of the log is reached, or `{error, Reason}` if an error occurs.

When `chunk/2,3` is used to wrap logs, the returned continuation may or may not be valid in the next call to `chunk`. This is because the log may wrap and delete the file into

which the continuation points. To make sure this does not happen, the log can be blocked during the search.

```
chunk_step(Log, Continuation, Step) -> {ok, Continuation2} | {error, Reason}
```

Types:

- Log = term()
- Continuation = start | cont()
- Step = int()
- Continuation2 = cont()
- Reason = end\_of\_log | term()

This function can be used in conjunction with `chunk/2,3` to search through a wrap log. It takes as argument a continuation as returned by `chunk/2,3` or `chunk_step/3`, and steps forward (or backward) `Step` files in the wrap log. The continuation returned points to the first log item in the new file.

If the wrap log is not full because all files have not been used yet, `{error, end_of_log}` is returned if trying to step outside the log.

```
close(Log) -> ok | {error, Reason}
```

This function closes a log file properly. This must be done before the system is stopped, or a log file with format `internal` is regarded as unclosed and the automatic repair procedure will be activated the next time the log is opened.

```
inc_wrap_log(Log) -> ok | {error, Reason}
```

This function forces the `disk_log` to start to log on the next log file. It can be used, for instance, if the user wants to read the latest logs and to be sure that it doesn't interfere with any possible new loggings. Another usage could be if the user wants to change to new a log file each day.

```
log(Log, Term) -> ok | {error, Reason}
```

```
blog(Log, Bytes) -> ok | {error, Reason}
```

Types:

- Log = term()
- Term = term()
- Bytes = binary() | [Byte]
- Byte = [Byte] | 0 =< integer() =< 255
- Reason = {file\_error, FileError} | term()

This function appends the term `Term` at the end of the log `Log`. It returns `ok` or `{error, Reason}` when the term has been written to disk. Terms are written by means of the normal `write()` function of the local operating system. Hence, there is no guarantee that the term has actually been written to the disk, it might linger in the operating system kernel for a while. To make sure the log is actually written to disk, the `sync/1` function must be called.

The `log` function is used for logs with format `internal`, and `blog` with format `external`.

If there is an error when writing to file, `{error, {file_error, FileError}}` is returned, with `FileError` as returned from `file:write`.

```
log_terms(Log, TermList) -> ok | {error, Reason}
blog_terms(Log, BytesList) -> ok | {error, Reason}
```

Types:

- Log = term()
- TermList = [term()]
- BytesList = [Bytes]
- Bytes = binary() | [Byte]
- Byte = [Byte] | 0 =< integer() =< 255

This function appends a list of items to the log. The difference between this function and the `log/2` function is that this function takes every term in `TermList` and produces a log item from it. This is *not* the same as logging a list of objects once! This function is more efficient than calling `log/2` for each item in the list.

The `log_terms` function is used for logs with format `internal`, and `blog_terms` with format `external`.

```
open(ArgL) -> OpenRet
```

Types:

- ArgL = [Opt]
- Opt = {name, term()} | {file, string()}, {linkto, LinkTo} | {repair, Repair} | {type, Type} | {format, Format} | {size, Size} | {distributed, [Node]} | {notify, bool()} | {head, Head} | {head\_func, {M,F,A}} | {mode, Mode}
- LinkTo = pid() | none
- Repair = true | false | truncate
- Type = halt | wrap
- Format = internal | external
- Size = integer() > 0 | infinity | {MaxBytes, MaxFiles}
- MaxBytes = integer() > 0
- MaxFiles = integer() > 0
- Rec = integer()
- Bad = integer()
- Node = atom()
- Head = none | term()
- Mode = read\_write | read\_only
- OpenRet = {ok, Name} | {repaired, Name, {recovered, Rec}, {badbytes, Bad}} | {error, {name\_already\_open, Name}} | {error, Reason} | DistRet
- DistRet = [{[Node, OpenRet]}, [{BadNode, {error, Reason}}]}

The `ArgL` parameter is a list of options which have the following meanings:

- {name, Name} specifies the name of the log. This is the name which must be passed as a parameter in all subsequent logging operations. A name must always be supplied.
- {file, Filename} specifies the name of the file which will be used to log terms. If this value is omitted and the name is either an atom or a string, the file name will default to `lists:concat([Name, ".LOG"])` for halt logs. For wrap logs, this will be the base name of the files. Each file in a wrap log will be called `<basename>.N`, where `N` is an integer. Each wrap log will also have two files called `<basename>.idx`, `<basename>.siz`.

- `{linkto, LinkTo}`. The log process can be set up to monitor a pid and then close the file properly if the pid should terminate. This pid is called the owner of the log. If the value `none` is supplied, the log file will remain open until explicitly closed. By default, the process which calls `open` owns the log.
- `{repair, Repair}`. If `false` is given, no automatic repair will be attempted. Instead, the tuple `{error, need_repair}` is returned if an attempt is made to open log file which was not properly closed. If `truncate` is given, the log file will be truncated, and thus create an empty log.
- `{type, Type}` is the type of the log. Default is `halt`. When a `halt` log reaches its maximum size, all attempts to log more items are rejected.
- `{format, Format}` specifies the format of the log items in the log. Default is `internal`. With `internal` format of the items in the log, the log file must be read with the function `chunk/2,3`. The format of the log file is internally defined, and it is not possible to view the file as ASCII text. With `external` format, however, all log items are written to the file exactly as they are, and it is the programmer's responsibility to format the items.
- `{size, Size}` specifies the size of the log. Default for `halt` logs is `infinity`. When wrap logs are used, the `Size` parameter is a 2-tuple `{MaxBytes, MaxFiles}`. The wrap log writes at most `MaxBytes` bytes on each file, it uses `MaxFiles` files before it wraps, and it truncates the first file.
- `{distributed, Nodes}`. This option should be used if a log should be globally visible and replicated on several nodes. If the log does not exist on any node, it is created. If the log does exist on some nodes, `Nodes` are joined to the existing nodes. The recommended way of using this functionality is to open the log with `{distributed, [node()]}` on each node. The module `pg2` is used to address the logs.
- `{notify, bool()}`. If `true`, the owner of the log is notified when certain events occur in the log. Default if `false`. The owner is sent one of the following messages when an event occurs:
  - `{disk_log, Node, Log, {wrap, NoLostItems}}` is sent when a wrap log has filled one of its files and a new file is opened.
  - `{disk_log, Node, Log, {truncated, NoLostItems}}` is sent when a log has been truncated, or dumped to a file.
  - `{disk_log, Node, Log, {read_only, Items}}` is sent when an asynchronous log attempt is made to a read only opened log file. `Items` is the items from the log attempt.
  - `{disk_log, Node, Log, full}` is sent when a `halt` log is full.
- `{head, Head}` specifies if a header should be written first on the log file. If the log is a wrap log, the `Head` is written first in each new file.
- `{head_func, {M,F,A}}` specifies that each time a new file is opened, `M:F(A)` is called. This function is supposed to return `{ok, Head}`. The `Head` is written first in each file. The `Head` should be a term if the format is `internal`, and a list of bytes (or a binary) otherwise.
- `{mode, Mode}` specifies if the log is opened in read only or read write mode. It defaults to `read_write`.

The `open/1` function returns `{ok, Name}` if the log file was successfully opened. If the file was successfully repaired, the tuple `{repaired, Name, {recovered, Rec}, {badbytes, Bad}}` is returned, where `Rec` is the number of whole Erlang terms found

in the file and `Bad` is the number of bytes in the file which were non-Erlang terms. If the `distributed` parameter was given to `open`, the function returns a list of successful replies and a list of erroneous replies. Each reply is tagged with the node name.

The `open/1` function ensures that the log server is started. Accordingly, it is not necessary to explicitly start the server first.

The function returns `{error, Reason}` for all other errors.

```
reopen(Log, File)
```

```
reopen(Log, File, Head)
```

```
breopen(Log, File, BHead) -> ok | {error, Reason}
```

Types:

- `Log = term()`
- `File = string()`
- `Head = term()`
- `BHead = binary() | [Byte]`
- `Byte = [Byte] | 0 =< integer() =< 255`

This function first renames the log file to `File` and then re-creates a new `Log` file. It is thus very efficient. If the `Head` or `BHead` arguments is given, this item is written first in the newly opened log file.

The `reopen/3` function is used for logs with format `internal`, and `breopen/3` with format `external`.

```
sync(Log) -> ok | {error, Reason}
```

Types:

- `Log = term()`

Ensures that the contents of the log is actually written to the disk. This is usually a pretty expensive operation.

```
truncate(Log)
```

```
truncate(Log, Head)
```

```
btruncate(Log, BHead) -> ok | {error, Reason}
```

Types:

- `Log = term()`
- `Head = term()`
- `BHead = binary() | [Byte]`
- `Byte = [Byte] | 0 =< integer() =< 255`

This function truncates a halt log. It cannot be used for wrap logs. If the `Head` or `BHead` arguments are given, this item is written first in the newly truncated log.

The `truncate/2` function is used for logs with format `internal`, and `btruncate/2` with format `external`.

```
unlock(Log) -> ok | {error, Reason}
```

Types:

- `Log = term()`

- Reason = not\_blocked | no\_such\_log

This function unblocks a log.

## See Also

file(3), pg2(3)

# erl\_boot\_server (Module)

This server is used to assist diskless Erlang systems which fetch all Erlang code from another machine.

This server is used to fetch all code, including the start script, if an Erlang system is started with the `-loader inet` command line flag. All hosts specified with the `-hosts Host` flag must have one instance of this server running.

This server can be started with the `kernel` configuration parameter `start_boot_server`.

## Exports

```
start(Slaves) -> {ok, Pid} | {error, What}
```

Types:

- Slaves = [Host]
- Host = atom()
- Pid = pid()
- What = void()

Starts the boot server. `Slaves` is a list of IP addresses for hosts which are allowed to use this server as a boot server.

```
start_link(Slaves) -> {ok, Pid} | {error, What}
```

Types:

- Slaves = [Host]
- Host = atom()
- Pid = pid()
- What = void()

Starts the boot server and links to the caller. This function is used to start the server if it is included in a supervision tree.

```
add_slave(Slave) -> ok | {error, What}
```

Types:

- Slave = Host
- Host = atom()
- What = void()

Adds a `Slave` node to the list of allowed slave hosts.

`delete_slave(Slave) -> ok | {error, What}`

Types:

- Slave = Host
- Host = atom()
- What = void()

Deletes a Slave node from the list of allowed slave hosts.

`which_slaves() -> Slaves`

Types:

- Slaves = [Host]
- Host = atom()

Returns the current list of allowed slave hosts.

## SEE ALSO

`init(3)`, `erl_prim_loader(3)`

# erl\_ddll (Module)

The `erl_ddll` module can load and link a linked-in driver, if run-time loading and linking of shared objects, or dynamic libraries, is supported by the underlying operating system.

## Exports

`start()` -> {ok, Pid} | {error, Reason}

Starts `ddll_server`. The error return values are the same as for `gen_server`.

`start_link()` -> {ok, Pid} | {error, Reason}

Starts `ddll_server` and links it to the calling process. The error return values are the same as for `gen_server`.

`stop()` -> ok

Stops `ddll_server`.

`load_driver(Path, Name)` -> ok | {error, ErrorDescriptor}

Types:

- Name = string() | atom()
- Path = string() | atom()

Loads and links the dynamic driver `Name`. `Name` must be sharable object/dynamic library. Two drivers with different `Paths` cannot be loaded under the same name. The number of dynamically loadable drivers are limited by the size of `driver_tab` in `config.c`.

If the server is not started the caller will crash.

`unload_driver(Name)` -> ok | {error, ErrorDescriptor}

Types:

- Name = string() | atom()

Unloads the dynamic driver `Name`. This will fail if any port programs are running the code that is being unloaded. Linked-in driver cannot be unloaded. The process must previously have called `load_driver/1` for the driver.

There is no guarantee that the memory where the driver was loaded is freed. This depends on the underlying operating system.

If the server is not started the caller will crash.

```
loaded_drivers() -> {ok, DriverList}
```

Types:

- DriverList = [Driver()]
- Driver = string()

Returns a list of all the available drivers, both (statically) linked-in and dynamically loaded ones.

If the server is not started the caller will crash.

```
format_error(ErrorDescriptor) -> string()
```

Takes an `ErrorDescriptor` which has been returned by one of `load_driver/2` and `unload_driver/1` and returns a string which describes the error or warning.

## Differences Between Statically Linked-in Drivers and Dynamically Loaded Drivers

Except for the following minor changes, all information in Appendix E of Concurrent Programming in Erlang, second edition, still applies.

The `driver_entry` struct has two new members: `finish` and `handle`.

Before the driver is unloaded, the `finish` function is called, without arguments, to give the driver writer a chance to clean up and release memory allocated in `driver_init`.

The member `handle` contains a pointer obtained from the operating system when the driver was loaded. Without this, the driver cannot be unloaded!

The `init` function in `struct driver_entry` is not used anymore. After the driver is loaded, the function `struct driver_entry *driver_init(void *)` is called with `handle` as argument. If the operating system loader cannot find a function called `driver_init`, the driver will not be loaded. The `driver_init` function *must* initialize a `struct driver_entry` and return a pointer to it.

Example:

```
#include <stdio.h>
#include "driver.h"
static long my_start();
static int my_stop(), my_read();
static struct driver_entry my_driver_entry;
/*
 * Initialize and return a driver entry struct
 */
struct driver_entry *driver_init(void *handle)
{
    my_driver_entry.init = null_func;      /* Not used */
    my_driver_entry.start = my_start;
    my_driver_entry.stop = my_stop;
    my_driver_entry.output = my_read;
    my_driver_entry.ready_input = null_func;
    my_driver_entry.ready_output = null_func;
    my_driver_entry.driver_name = "my_driver";
}
```

```
my_driver_entry.finish = null_func;
my_driver_entry.handle = handle; /* MUST set this!!! */
return &my_driver_entry;
}
```

## config.c

The size of the `driver_tab` array, defined in `config.c`, limits the number of dynamically loadable drivers.

## Compiling Your Driver

Please refer to your C compiler or operating system documentation for information about producing a sharable object or DLL.

The include file `driver.h` is found in the `usr/include` directory of the Erlang installation.

# erl\_prim\_loader (Module)

The `erl_prim_loader` is used to load all Erlang modules into the system. The start script is also fetched with the low level loader.

The `erl_prim_loader` knows about the environment and how to fetch modules. The loader could, for example, fetch files using the file system (with absolute file names as input), or a database (where the binary format of a module is stored).

The `-loader Loader` command line flag can be used to choose the method used by the `erl_prim_loader`. Two Loader methods are supported by the Erlang system: `efile` and `inet`. If another loader is required, then it has to be implemented by the user. The Loader provided by the user must fulfill the protocol defined below, and it is started with the `erl_prim_loader` by evaluating `open_port({spawn, Loader}, [binary])`.

## Exports

```
start(Id,Loader,Hosts) -> {ok, Pid} | {error, What}
```

Types:

- Id = term()
- Loader = atom()
- Hosts = [Host]
- Host = atom()
- Pid = pid()
- What = term()

Starts the Erlang low level loader. This function is called by the `init` process (and module). The `init` process reads the command line flags `-id Id`, `-loader Loader`, and `-hosts Hosts`. These are the arguments supplied to the `start/3` function.

If `-loader` is not given, the default loader is `efile` which tells the system to read from the file system.

If `-loader` is `inet`, the `-id Id`, `-hosts Hosts`, and `-setcookie Cookie` flags must also be supplied. `Hosts` identifies hosts which this node can contact in order to load modules. One Erlang system with a `erl_boot_server` process must be started on each of hosts given in `Hosts` in order to answer the requests. See `erl_boot_server(3)`.

If `-loader` is something else, the given port program is started. The port program is supposed to follow the protocol specified below.

```
get_file(File) -> {ok, Bin, FullName} | error
```

Types:

- File = string()

- Bin = binary()
- FullName = string()

This function fetches a file using the low level loader. `File` is either an absolute file name or just the name of the file, for example "lists.jam". If an internal path is set to the loader, this path is used to find the file. If a user supplied loader is used, the path can be stripped off if it is obsolete, and the loader does not use a path. `FullName` is the complete name of the fetched file. `Bin` is the contents of the file as a binary.

`get_path()` -> {ok, Path}

Types:

- Path = [Dir]
- Dir = string()

This function gets the path set in the loader. The path is set by the `init` process according to information found in the start script.

`set_path(Path)` -> ok

Types:

- Path = [Dir]
- Dir = string()

This function sets the path of the loader if `init` interprets a `path` command in the start script.

## Protocol

The following protocol must be followed if a user provided loader port program is used. The Loader port program is started with the command `open_port({spawn, Loader}, [binary])`. The protocol is as follows:

Function	Send	Receive
<code>get_file</code>	[102   FileName]	[121   BinaryFile] (on success) [122] (failure)
<code>stop</code>	<code>eof</code>	<code>terminate</code>

## Command Line Flags

The `erl_prim_loader` module interprets the following flags:

**-loader *Loader*** Specifies the name of the loader used by `erl_prim_loader`. *Loader* can be `efile` (use the local file system), or `inet` (load using the `boot_server` on another Erlang node). If *Loader* is user defined, the defined *Loader* port program is started.

If the `-loader` flag is omitted, it defaults to `efile`.

**-hosts *Hosts*** Specifies which other Erlang nodes the `inet` loader can use. This flag is mandatory if the `-loader inet` flag is present. On each host, there must be on Erlang node with the `erl_boot_server` which handles the load requests. *Hosts* is a list of IP addresses (hostnames are not acceptable).

**-id *Id*** Specifies the identity of the Erlang system. If the system runs as a distributed node, *Id* must be identical to the name supplied with the `-sname` or `-name` distribution flags.

**-setcookie *Cookie*** Specifies the cookie of the Erlang system. This flag is mandatory if the `-loader inet` flag is present.

## SEE ALSO

`init(3)`, `erl_boot_server(3)`

## erlang (Module)

By convention, Built In Functions (BIFs) are seen as being in the module `erlang`. Thus, both the calls `atom_to_list(Erlang)` and `erlang:atom_to_list(Erlang)` are identical.

BIFs may fail for a variety of reasons. All BIFs fail if they are called with arguments of an incorrect type. For example, `atom_to_list/1` will fail if it is called with an argument which is not an atom. If this type of failure is not within the scope of a catch (and the BIF is not called within a guard; see below), it will cause the process to make the call to exit, and an EXIT signal with the associated reason `badarg` will be sent to all linked processes. The other reasons that may make BIFs fail are described in connection with the description of each individual BIF.

A few BIFs may be used in guard tests, for example:

```
tuple_5(Something) when size(Something) == 5 ->
    is_tuple_size_5;
tuple_5(_) ->
    is_something_else.
```

Here the BIF `size/1` is used in a guard. If `size/1` is called with a tuple, it will return the size of the tuple (i.e., how many elements there are in the tuple). In the above example, `size/1` is used in a guard which tests if its argument `Something` is a tuple and, if it is, whether it is of size 5. In this case, calling `size` with an argument other than a tuple will cause the guard to fail and execution will continue with the next clause. Suppose `tuple_5/1` is written as follows:

```
tuple_5(Something) ->
    case size(Something) of
        5 -> is_tuple_size_5;
        _ -> is_something_else
    end.
```

In this case, `size/1` is not in a guard. If `Something` is not a tuple, `size/1` will fail and cause the process to exit with the associated reason `badarg` (see above).

Some of the BIFs in this chapter are optional in Erlang implementations, and not all implementations will include them.

The following descriptions indicate which BIFs can be used in guards and which BIFs are optional.

## Exports

`abs(Number)`

Returns an integer or float which is the arithmetical absolute value of the argument `Number` (integer or float).

```
> abs(-3.33).  
3.33000  
> abs(-3).  
3
```

This BIF is allowed in guard tests.

Failure: `badarg` if the argument is not an integer or a float.

`apply({Module, Function}, ArgumentList)`

This is equivalent to `apply(Module, Function, ArgumentList)`.

`apply(Module, Function, ArgumentList)`

Returns the result of applying `Function` in `Module` on `ArgumentList`. The applied function must have been exported from `Module`. The arity of the function is the length of `ArgumentList`.

```
> apply(lists, reverse, [[a, b, c]]).  
[c,b,a]
```

`apply` can be used to evaluate BIFs by using the module name `erlang`.

```
> apply(erlang, atom_to_list, ['Erlang']).  
"Erlang"
```

Failure: `error_handler:undefined_function/3` is called if `Module` has not exported `Function/Arity`. The error handler can be redefined (see the BIF `process_flag/2`). If the `error_handler` is undefined, or if the user has redefined the default `error_handler` so the replacement module is undefined, an error with the reason `undef` will be generated.

`atom_to_list(Atom)`

Returns a list of integers (Latin-1 codes), which corresponds to the text representation of the argument `Atom`.

```
> atom_to_list('Erlang').  
"Erlang"
```

Failure: `badarg` if the argument is not an atom.

`binary_to_list(Binary)`

Returns a list of integers which correspond to the bytes of `Binary`.

`binary_to_list(Binary, Start, Stop)`

As `binary_to_list/1`, but it only returns the list from position `Start` to position `Stop`. `Start` and `Stop` are integers. Positions in the binary are numbered starting from 1.

`binary_to_term(Binary)`

Returns an Erlang term which is the result of decoding the binary `Binary`. `Binary` is encoded in the Erlang external binary representation. See `term_to_binary/1`.

`erlang:cancel_timer(Ref)`

`cancel_timer(Ref)` cancels a timer, where `Ref` was returned by either `send_after/3` or `start_timer/3`. If the timer was there to be removed, `cancel_timer/1` returns the time in ms left until the timer would have expired, otherwise `false` (which may mean that `Ref` was never a timer, or that it had already been cancelled, or that it had already delivered its message).

Note: usually, cancelling a timer doesn't guarantee that the message hasn't already been delivered to the message queue. However, in the special case of a process `P` cancelling a timer which would have sent a message to `P` itself, attempting to read the timeout message from the queue is guaranteed to remove the timeout in that situation:

```
cancel_timer(Ref),
receive
    {timeout, Ref, _} ->
        ok
after 0 ->
    ok
end
```

Failure: `badarg` if `Ref` isn't a reference.

`erlang:check_process_code(Pid, Module)`

Returns `true` if the process `Pid` is executing an old version of `Module`, if the current call of the process executes code for an old version of the module, or if the process has references to an old version of the module. Otherwise, it returns `false`.

```
> erlang:check_process_code(Pid, lists).
false
```

This is an optional BIF.

Failure: `badarg`, if the argument is not a `Pid`.

`concat_binary(ListOfBinaries)`

Concatenates a list of binaries `ListOfBinaries` into one binary.

`date()`

Returns the current date as `{Year, Month, Day}`

```
> date().
{1995, 2, 19}
```

`erlang:delete_module(Module)`

Moves the current version of the code of `Module` to the old version and deletes all export references of `Module`. Returns `undefined` if the module does not exist, otherwise `true`.

```
> delete_module(test).
true
```

This is an optional BIF.

Failure: `badarg` if there is already an old version of the module (see BIF `purge_module/1`).

### Warning:

In normal Erlang implementations code handling - which includes loading, deleting, and replacing modules - is performed in the module code. This BIF is intended for use with the implementation of the module code and should not be used elsewhere.

`erlang:demonitor(Ref)`

If `Ref` is a reference which the current process obtained by calling `erlang:monitor/2`, the monitoring is turned off.

It is an error if `Ref` refers to a monitoring started by another process. Not all such cases are cheap to check; if checking is cheap, the call fails with `badarg` (for example if `Ref` is a remote reference).

`erlang:disconnect_node(Node)`

Forces the disconnection of a node. This will appear to the node `Node` as if the current node has crashed. This BIF is mainly used in the Erlang network authentication protocols. Returns `true` if disconnection succeeds, otherwise `false`.

Failure: `badarg` if `Node` is not an atom.

`element(N, Tuple)`

Returns the `N`th element (numbering from 1) of `Tuple`.

```
> element(2, {a, b, c}).
b
```

Failure: `badarg` if  $N < 1$ , or  $N > \text{size}(\text{Tuple})$ , or if the argument `Tuple` is not a tuple. Allowed in guard tests.

`erase()`

Returns the process dictionary and deletes it.

```
> put(key1, {1, 2, 3}), put(key2, [a, b, c]), erase().
[{key1, {1, 2, 3}}, {key2, [a, b, c]}]
```

`erase(Key)`

Returns the value associated with `Key` and deletes it from the process dictionary.

Returns `undefined` if no value is associated with `Key`. `Key` can be any Erlang term.

```
> put(key1, {merry, lambs, are, playing}),
  X = erase(key1), {X, erase(key1)}.
{{merry, lambs, are, playing}, undefined}
```

`exit(Reason)`

Stops the execution of the current process with the reason `Reason`. Can be caught. `Reason` is any Erlang term. Since evaluating this function causes the process to terminate, it has no return value.

```
> exit(foobar).
** exited: foobar **
> catch exit(foobar).
{'EXIT', foobar}
```

`exit(Pid, Reason)`

Sends an EXIT message to the process `Pid`. Returns `true`.

```
> exit(Pid, goodbye).
true
```

**Note:**

The above is not necessarily the same as:

```
Pid ! {'EXIT', self(), goodbye}
```

The above two alternatives are the same if the process with the process identity `Pid` is trapping exits. However, if `Pid` is not trapping exits, the `Pid` itself will exit and propagate EXIT signals in turn to its linked processes.

If the reason is the atom `kill`, for example `exit(Pid, kill)`, an untrappable EXIT signal will be sent to the process `Pid`. In other words, the process `Pid` will be unconditionally killed.

Returns `true`.

Failure: `badarg` if `Pid` is not a `Pid`.

`float(Number)`

Returns a float by converting `Number` to a float.

```
> float(55).
55.0000
```

**Note:**

`float/1` is allowed in guard tests, but it tests whether the argument is a float or not.

```
-module(t).

f(F) when float(F) -> float;
f(F) -> not_a_float.

1> t:f(1).
not_a_float
2> t:f(1.0).
float
3>
```

Failure: badarg if the argument is not a float or an integer.

`float_to_list(Float)`

Returns a list of integers (ASCII codes) which corresponds to `Float`.

```
> float_to_list(7.0).
"7.0000000000000000000000e+00"
```

Failure: badarg if the argument is not a float.

`get()`

Returns the process dictionary as a list of `{Key, Value}` tuples.

```
> put(key1, merry), put(key2, lambs),
   put(key3, {are, playing}), get().
[{key1, merry}, {key2, lambs}, {key3, {are, playing}}]
```

`get(Key)`

Returns the value associated with `Key` in the process dictionary, and undefined if no value is associated with `Key`. `Key` can be any Erlang term.

```
> put(key1, merry), put(key2, lambs),
   put({any, [valid, term]}, {are, playing}),
   get({any, [valid, term]}).
{are, playing}
```

`erlang:get_cookie()`

Returns the “magic cookie” of the current node, if the node is alive; otherwise the atom `nocookie`.

`get_keys(Value)`

Returns a list of keys which corresponds to `Value` in the process dictionary.

```
> put(mary, {1, 2}), put(had, {1, 2}), put(a, {1, 2}),
   put(little, {1, 2}), put(dog, {1, 3}), put(lamb, {1, 2}),
   get_keys({1, 2}).
[mary, had, a, little, lamb]
```

### group\_leader()

Every process is a member of some process group and all groups have a leader.

This BIF returns the process identifier `Pid` of the group leader for the process which evaluates the BIF. When a process is spawned, the group leader of the spawned process is the same as that of the process which spawned it. Initially, at system start-up, `init` is both its own group leader and the group leader of all processes.

### group\_leader(Leader, Pid)

Sets the group leader of `Pid` to `Leader`. Typically, this is used when a processes started from a certain shell should have another group leader than `init`. The process `Leader` is normally a process with an I/O protocol. All I/O from this group of processes are thus channeled to the same place.

### halt()

Halts the Erlang system and indicates normal exit to the calling environment. Has no return value.

```
> halt().
unix_prompt%
```

### halt(Status)

`Status` must be a non-negative integer, or a string. Halts the Erlang system. Has no return value. If `Status` is an integer, it is returned as an exit status of Erlang to the calling environment. If `Status` is a string, produces an Erlang crash dump with `String` as slogan, and then exits with a non-zero status code.

Note that on many platforms, only the status codes 0-255 are supported by the operating system.

### erlang:hash(Term, Range)

Returns a hash value for `Term` within the range `1..Range`.

### hd(List)

Returns the first item of `List`.

```
> hd([1,2,3,4,5]).
1
```

Allowed in guard tests.

Failure: `badarg` if `List` is the empty list `[]`, or is not a list.

### erlang:info(What)

This BIF is optional and may be removed or changed in future releases of Erlang. What can be any of the atoms `info`, `procs`, `loaded`, or `dist`. The BIF returns information of the different 'topics' as binary data objects.

Failure: `badarg` if What is not one of the atoms shown above.

`integer_to_list(Integer)`

Returns a list of integers (ASCII codes) which correspond to `Integer`.

```
> integer_to_list(77).  
"77"
```

Failure: `badarg` if the argument is not an integer.

`is_alive()`

Returns the atom `true` if the current node is alive; i.e., if the node can be part of a distributed system. Otherwise, it returns the atom `false`.

`is_process_alive(Pid)`

`Pid` must refer to a process on the current node. Returns the atom `true` if the process is alive, i.e., has not exited. Otherwise, it returns the atom `false`. This is the preferred way to check whether a process exists. Unlike `process_info/[1,2]`, `is_process_alive/1` does not report zombie processes as alive.

`length(List)`

Returns the length of `List`.

```
> length([1,2,3,4,5,6,7,8,9]).  
9
```

Allowed in guard tests.

Failure: `badarg` if the argument is not a proper list.

`link(Pid)`

Creates a link to the process (or port) `Pid`, if there is not such a link already. If a process attempts to create a link to itself, nothing is done. Returns `true`.

Failure: `badarg` if the argument is not a `Pid` or port. Sends the EXIT signal `noproc` to the process which evaluates `link` if the argument is the `Pid` of a process which does not exist.

`list_to_atom(CharIntegerList)`

Returns an atom whose text representation is the integers (Latin-1 codes) in `CharIntegerList`.

```
> list_to_atom([69, 114, 108, 97, 110, 103]).  
'Erlang'
```

Failure: `badarg` if the argument is not a list of integers, or if any integer in the list is not an integer in the range [0, 255].

`list_to_binary(ListOfIntegers)`

Returns a binary which is made from the integer list `ListOfIntegers`.

Failure: `badarg` if the argument is not a list of integers, or if any integer in the list is not an integer in the range `[0, 255]`.

`list_to_float(AsciiIntegerList)`

Returns a float whose text representation is the integers (ASCII-values) in `AsciiIntegerList`.

```
> list_to_float([50,46,50,48,49,55,55,54,52,101,43,48]).
2.20178
```

Failure: `badarg` if the argument is not a list of integers, or if `AsciiIntegerList` contains a bad representation of a float.

`list_to_integer(AsciiIntegerList)`

Returns an integer whose text representation is the integers (ASCII-values) in `AsciiIntegerList`.

```
> list_to_integer([49, 50, 51]).
123
```

Failure: `badarg` if the argument is not a list of integers, or if `AsciiIntegerList` contains a bad representation of an integer.

`list_to_pid(AsciiIntegerList)`

Returns a `Pid` whose text representation is the integers (ASCII-values) in `AsciiIntegerList`. This BIF is intended for debugging, and in the Erlang operating system. *It should not be used in application programs.*

```
> list_to_pid("<0.4.1>").
<0.4.1>
```

Failure: `badarg` if the argument is not a list of integers, or `AsciiIntegerList` contains a bad representation of a `Pid`.

`list_to_tuple(List)`

Returns a tuple which corresponds to `List`. `List` can contain any Erlang terms.

```
> list_to_tuple([mary, had, a, little, {dog, cat, lamb}]).
{mary, had, a, little, {dog, cat, lamb}}
```

Failure: `badarg` if `List` is not a proper list.

`erlang:load_module(Module, Binary)`

If `Binary` contains the object code for the module `Module`, this BIF loads that object code. Also, if the code for the module `Module` already exists, all export references are replaced so they point to the newly loaded code. The previously loaded code is kept in the system as 'old code', as there may still be processes which are executing that code. It returns either `{module, Module}`, where `Module` is the name of the module which has been loaded, or `{error, Reason}` if load fails. `Reason` is one of the following:

`badfile` If the object code in `Binary` has an incorrect format.

`not_purged` If `Binary` contains a module which cannot be loaded because old code for this module already exists (see the BIFs `purge_module` and `delete_module`).

`badfile` If the object code contains code for another module than `Module`

**Warning:**

Code handling - which includes loading, deleting, and replacing of modules - is done by the module `code` in normal Erlang implementations. This BIF is intended for the implementation of the module named `code` and should not be used elsewhere.

This is an optional BIF.

Failure: `badarg` if the first argument is not an atom, or the second argument is not a binary.

`erlang:localtime()`

Returns the current local date and time `{{Year, Month, Day}, {Hour, Minute, Second}}`.

The time zone and daylight saving time correction depend on the underlying OS.

```
> erlang:localtime().
{{1996,11,6},{14,45,17}}
```

`erlang:localtime_to_universaltime(DateTime)`

Converts local date and time in `DateTime` to Universal Time Coordinated (UTC), if this is supported by the underlying OS. Otherwise, no conversion is done and `DateTime` is returned. The return value is of the form `{{Year, Month, Day}, {Hour, Minute, Second}}`.

Failure: `badarg` if the argument is not a valid date and time tuple `{Year, Month, Day}, {Hour, Minute, Second}`.

```
> erlang:localtime_to_universaltime({{1996,11,6},{14,45,17}}).
{{1996,11,6},{13,45,17}}
```

`make_ref()`

Returns an almost unique reference.

The returned reference will reoccur after approximately 260 000 calls; hence it is not unique. This is an error that will be corrected in later releases.

In order to obtain a true unique reference use a combination of calls to `erlang:now/0` and `make_ref/0` as follows:

```
{erlang:now(), make_ref() }
```

```
> make_ref().
#Ref
```

`erlang:module_loaded(Module)`

Returns the atom `true` if the module contained in atom `Module` is loaded, otherwise it returns the atom `false`. It does not attempt to load the module.

**Warning:**

This BIF is intended for the implementation of the module named `code` and should not be used anywhere else. Use `code:is_loaded/1` instead.

```
> erlang:module_loaded(lists).
true
```

This is an optional BIF.

Failure: `badarg` if the argument is not an atom.

`erlang:monitor(Type, Item)`

The current process starts monitoring `Item`, and will be notified when `Item` dies, with a message `{'DOWN', Ref, Type, Item, Info}`, where `Ref` is the value returned by the call to `erlang:monitor/2`, and `Info` gives additional information. The message is also sent if `Item` is already dead. The monitored item must be local. The value returned can be used for disabling the monitor (see `erlang:demonitor/1`). The currently allowed arguments are the atom `process` for `Type`, and a `pid` for `Item`; `Info` in the message is the exit reason of the process (or `noproc`, as for `link/1`).

Making several calls to `erlang:monitor/2` for the same item is not an error; it results in several completely independent monitorings.

`monitor_node(Node, Flag)`

Monitors the status of the node `Node`. If `Flag` is `true`, monitoring is turned on; if `Flag` is `false`, monitoring is turned off. Calls to the BIF are accumulated. This is shown in the following example, where a process is already monitoring the node `Node` and a library function is called:

```
monitor_node(Node, true),
    ... some operations
monitor_node(Node, false),
```

After the call, the process is still monitoring the node.

If `Node` fails or does not exist, the message `{nodedown, Node}` is delivered to the process. If a process has made two calls to `monitor_node(Node, true)` and `Node` terminates, two `nodedown` messages are delivered to the process. If there is no connection to `Node`, there will be an attempt to create one. If this fails, a `nodedown` message is delivered.

Returns `true`.

Failure: `badarg` if `Flag` is not `true` or `false`, and `badarg` if `Node` is not an atom indicating a remote node, or if the local node isn't alive.

`node()`

Returns the name of the current node. If it is not a networked node but a local Erlang system, the atom `nonode@nohost` is returned.

Allowed in guard tests.

`node(Arg)`

Returns the node where `Arg` is located. `Arg` can be a `Pid`, a reference, or a port.

Allowed in guard tests.

Failure: `badarg` if `Arg` is not a `Pid`, reference, or port.

`nodes()`

Returns a list of all known nodes in the system, excluding the current node.

`now()`

Returns the tuple `{MegaSecs, Secs, Microsecs}`

which is the elapsed time since 00:00 GMT, January 1, 1970 (zero hour) on the assumption that the underlying OS supports this. Otherwise, some other point in time is chosen. It is also guaranteed that subsequent calls to this BIF returns continuously increasing values. Hence, the return value from `now()` can be used to generate unique time-stamps. It can only be used to check the local time of day if the time-zone info of the underlying operating system is properly configured.

`open_port(PortName, PortSettings)`

Returns a port identifier as the result of opening a new Erlang port. A port can be seen as an external Erlang process. `PortName` is one of the following:

`{spawn, Command}` Starts an external program. `Command` is the name of the external program which will be run. `Command` runs outside the Erlang work space unless an Erlang driver with the name `Command` is found. If found, that driver will be started. A driver runs in the Erlang workspace, which means that it is linked with the Erlang runtime system.

When starting external programs on Solaris, the system call `vfork` is used in preference to `fork` for performance reasons, although it has a history of being less robust. If there are problems with using `vfork`, setting the environment variable `ERL_NO_VFORK` to any value will cause `fork` to be used instead.

`Atom` *This use of `open_port()` is obsolete and will be removed in a future version of Erlang. Use the `file` module instead.* The atom is assumed to be the name of an external resource. A transparent connection is established between Erlang and the resource named by the atom `Atom`. The characteristics of the port depend on the type of resource. If `Atom` represents a normal file, the entire contents of the file is sent to the Erlang process as one or more messages. When messages are sent to the port, it causes data to be written to the file.

`{fd, In, Out}` Allows an Erlang process to access any currently opened file descriptors used by Erlang. The file descriptor `In` can be used for standard input, and the file descriptor `Out` for standard output. It is only used for various servers in the Erlang operating system (`shell` and `user`). Hence, its use is very limited.

`PortSettings` is a list of settings for the port. Valid values are:

- `{packet, N}` Messages are preceded by their length, sent in `N` bytes, with the most significant byte first. Valid values for `N` are 1, 2, or 4.
- `stream` Output messages are sent without packet lengths. A user-defined protocol must be used between the Erlang process and the external object.
- `{line, N}` Messages are delivered on a per line basis. Each line (delimited by the OS-dependent newline sequence) is delivered in one single message. The message data format is `{Flag, Line}`, where `Flag` is either `eol` or `noeol` and `Line` is the actual data delivered (without the newline sequence).
- `N` specifies the maximum line length in bytes. Lines longer than this will be delivered in more than one message, with the `Flag` set to `noeol` for all but the last message. If end of file is encountered anywhere else than immediately following a newline sequence, the last line will also be delivered with the `Flag` set to `noeol`. In all other cases, lines are delivered with `Flag` set to `eol`.
- The `{packet, N}` and `{line, N}` settings are mutually exclusive.
- `{cd, Dir}` This is only valid for `{spawn, Command}`. The external program starts using `Dir` as its working directory. `Dir` must be a string. Not available on VxWorks.
- `{env, Environment}` This is only valid for `{spawn, Command}`. The environment of the started process is extended using the environment specifications in `Environment`. `Environment` should be a list of tuples `{Name, Value}`, where `Name` is the name of an environment variable, and `Value` is the value it is to have in the spawned port process. Both `Name` and `Value` must be strings. The one exception is `Value` being the atom `false` (in analogy with `os:getenv/1`), which removes the environment variable. Not available on VxWorks.
- `exit_status` This is only valid for `{spawn, Command}` where `Command` refers to an external program. When the external process connected to the port exits, a message of the form `{Port, {exit_status, Status}}` is sent to the connected process, where `Status` is the exit status of the external process. If the program aborts, on Unix the same convention is used as the shells do (i.e. `128+signal`). If the `eof` option has been given as well, the `eof` message and the `exit_status` message appear in an unspecified order. If the port program closes its `stdout` without exiting, the `exit_status` option won't work.
- `use_stdio` This is only valid for `{spawn, Command}`. It allows the standard input and output (file descriptors 0 and 1) of the spawned (UNIX) process for communication with Erlang.
- `nouse_stdio` The opposite of the above. Uses file descriptors 3 and 4 for communication with Erlang.
- `stderr_to_stdout` Affects ports to external programs. The executed program gets its standard error file redirected to its standard output file. `stderr_to_stdout` and `nouse_stdio` are mutually exclusive.
- `in` The port can only be used for input.
- `out` The port can only be used for output.
- `binary` All I/O from the port are binary data objects as opposed to lists of bytes.
- `eof` The port will not be closed at the end of the file and produce an EXIT signal. Instead, it will remain open and a `{Port, eof}` message will be sent to the process holding the port.

The default is `stream` for all types of port and `use_stdio` for spawned ports.

Failure: `badarg` if the format of `PortName` or `PortSettings` is incorrect. If the port cannot be opened, the exit reason is the Posix error code which most closely describes

the error, or `EINVAL` if no Posix code is appropriate. The following Posix error codes may appear:

`ENOMEM` There was not enough memory to create the port.

`EAGAIN` There are no more available operating system processes.

`ENAMETOOLONG` The external command given was too long.

`EMFILE` There are no more available file descriptors.

`ENFILE` A file or port table is full.

During use of a port opened using `{spawn, Name}`, errors arising when sending messages to it are reported to the owning process using exit signals of the form `{'EXIT', Port, PosixCode}`. Posix codes are listed in the documentation for the `file` module.

`pid_to_list(Pid)`

Returns a list which corresponds to the process `Pid`.

**Warning:**

This BIF is intended for debugging and for use in the Erlang operating system. It should not be used in application programs.

```
> pid_to_list(whereis(init)).
"<0.0.0>"
```

Failure: `badarg` if the argument is not a `Pid`.

`erlang:port_info(Port, Item)`

Returns information about the port `Port` as specified by `Item`, which can be any one of the atoms `id`, `connected`, `links`, `name`, `input`, or `output`.

`{id, Index}` `Index` is the internal index of the port. This index may be used to separate ports.

`{connected, Pid}` `Pid` is the process connected to the port.

`{links, ListOfPids}` `ListOfPids` is a list of `Pids` with processes to which the port has a link.

`{name, String}` `String` is the command name set by `open_port`.

`{input, Bytes}` `Bytes` is the total number of bytes read from the port.

`{output, Bytes}` `Bytes` is the total number of bytes written to the port.

All implementations may not support all of the above `Items`. Returns `undefined` if the port does not exist.

Failure: `badarg` if `Port` is not a process identifier, or if `Port` is a port identifier of a remote process.

`erlang:ports()`

Returns a list of all ports on the current node.

`erlang:pre_loaded()`

Returns a list of Erlang modules which are pre-loaded in the system. As all loading of code is done through the file system, the file system must have been loaded previously. Hence, at least the module `init` must be pre-loaded.

`erlang:process_display(Pid, Type)`

Writes information about the local process `Pid` on standard error. The currently allowed value for the atom `Type` is `backtrace`, which shows the contents of the stack, including information about the call chain, with the most recent data printed last. The format of the output is not further defined. `Pid` may be a zombie process.

`process_flag(Flag, Option)`

Sets certain flags for the process which calls this function. Returns the old value of the flag.

`process_flag(trap_exit, Boolean)` When `trap_exit` is set to `true`, EXIT signals arriving to a process are converted to `{'EXIT', From, Reason}` messages, which can be received as ordinary messages. If `trap_exit` is set to `false`, the process exits if it receives an EXIT signal other than `normal` and the EXIT signal is propagated to its linked processes. Application processes should normally not trap exits.

`process_flag(error_handler, Module)` This is used by a process to redefine the error handler for undefined function calls and undefined registered processes. Inexperienced users should not use this flag since code autoloading is dependent on the correct operation of the error handling module.

`process_flag(priority, Level)` This sets the process priority. `Level` is an atom. All implementations support three priority levels, `low`, `normal`, and `high`. The default is `normal`.

Failure: `badarg` if `Flag` is not an atom, or is not a recognized flag value, or if `Option` is not a recognized term for `Flag`.

`process_info(Pid)`

Returns a long list which contains information about the process `Pid`. This BIF is only intended for debugging. It should not be used for any other purpose. The list returned contains the following tuples. The order in which these tuples are returned is not defined, nor are all the tuples mandatory.

`{current_function, {Module, Function, Arguments}}` `Module, Function, Arguments` is the current function call of the process.

`{dictionary, Dictionary}` `Dictionary` is the dictionary of the process.

`{error_handler, Module}` `Module` is the error handler module used by the process (for undefined function calls, for example).

`{group_leader, Groupleader}` `Groupleader` is group leader for the I/O of the process.

`{heap_size, Size}` `Size` is the heap size of the process in heap words.

`{initial_call, {Module, Function, Arity}}` `Module, Function, Arity` is the initial function call with which the process was spawned.

`{links, ListOfPids}` `ListOfPids` is a list of `Pids`, with processes to which the process has a link.

`{message_queue_len, MessageQueueLen}` `MessageQueueLen` is the number of messages currently in the message queue of the process. This is the length of the list `MessageQueue` returned as the info item `messages` (see below).

`{messages, MessageQueue}` `MessageQueue` is a list of the messages to the process, which have not yet been processed.

`{priority, Level}` `Level` is the current priority level for the process. Only `low` and `normal` are always supported.

`{reductions, Number}` `Number` is the number of reductions executed by the process.

`{registered_name, Atom}` `Atom` is the registered name of the process. If the process has no registered name, this tuple is not present in the list.

`{stack_size, Size}` `Size` is the stack size of the process in stack words.

`{status, Status}` `Status` is the status of the process. `Status` is `waiting` (waiting for a message), `running`, `runnable` (ready to run, but another process is running), `suspended` (suspended on a “busy” port or by the `trace/3` BIF), or `exiting` (if the process has exited, but remains as a zombie).

`{trap_exit, Boolean}` `Boolean` is `true` if the process is trapping exits, otherwise it is `false`.

Failure: `badarg` if the argument is not a `Pid`, or if `Pid` is a `Pid` of a remote process.

`process_info(Pid, Item)`

Returns information about the process `Pid` as specified by `Item`, in the form `{Item, Info}`. `Item` can be any one of the atoms `current_function`, `dictionary`, `error_handler`, `group_leader`, `heap_size`, `initial_call`, `links`, `memory`, `message_queue_len`, `messages`, `priority`, `reductions`, `registered_name`, `stack_size`, `status`, `trap_exit` or `exit`.

Returns `undefined` if no information is known about the process.

`process_info` can be used to obtain information about processes which have exited but whose data are still kept, so called zombie processes. To determine whether to keep information about dead processes, use the BIF `erlang:system_flag/2`. Since `process_info` does not necessarily return `undefined` for a dead process, use `is_process_alive/1` to check whether a process is alive.

Item `exit` returns `[]` if the process is alive, or `{exit, Reason}` if the process has exited, where `Reason` is the exit reason.

Item `registered_name` returns `[]` if the process has no registered name. If the process is a zombie, the registered name it had when it died is returned.

Item `memory` returns `{memory, Size}`, where `Size` is the size of the process in bytes. This includes stack, heap and internal structures.

Not all implementations support every one of the above `Items`.

Failure: `badarg` if `Pid` is not a process identifier, or if `Pid` is a process identifier of a remote process.

`processes()`

Returns a list of all processes on the current node, including zombie processes. See `system_flag/2` [page 82].

```
> processes().
[<0.0.1>, <0.1.1>, <0.2.1>, <0.3.1>, <0.4.1>, <0.6.1>]
```

`erlang:purge_module(Module)`

Removes old code for `Module`. Before this BIF is used, `erlang:check_process_code/2` should be called to check that no processes are executing old code in this module.

**Warning:**

In normal Erlang implementations, code handling - which is loading, deleting and replacing modules - is evaluated by the module code. This BIF is intended to be used by the implementation of the module code and should not be used in any other place.

This is an optional BIF.

Failure: `badarg` if `Module` does not exist.

`put(Key, Value)`

Adds a new `Value` to the process dictionary and associates it with `Key`. If a value is already associated with `Key`, that value is deleted and replaced by the new value `Value`. It returns any value previously associated with `Key`, or `undefined` if no value was associated with `Key`. `Key` and `Value` can be any valid Erlang terms.

**Note:**

The values stored when `put` is evaluated within the scope of a `catch` will not be retracted if a `throw` is evaluated, or if an error occurs.

```
> X = put(name, walrus), Y = put(name, carpenter),
  Z = get(name),
  {X, Y, Z}.
{undefined, walrus, carpenter}
```

`register(Name, Pid)`

Associates the name `Name` with the process identity `Pid`. `Name`, which must be an atom, can be used instead of a `pid` in the send operator (`Name ! Message`).

Returns `true`.

Failure: `badarg` if `Pid` is not an active process, or if `Pid` is a process on another node, or if the name `Name` is already in use, or if the process is already registered (it already has a name), or if the name `Name` is not an atom, or if `Name` is the atom `undefined`.

`registered()`

Returns a list of names which have been registered using `register/2`.

```
> registered().
[code_server, file_server, init, user, my_db]
```

`round(Number)`

Returns an integer by rounding the number `Number`. Allowed in guard tests.

```
> round(5.5).  
6
```

Failure: `badarg` if the argument is not a float (or an integer).

`self()`

Returns the process identity of the calling process. Allowed in guard tests.

```
> self().  
<0.16.1>
```

`erlang:send_after(Time, Pid, Msg)`

`Time` is a non-negative integer, `Pid` is either a pid or an atom, and `Msg` is any Erlang term. The function returns a reference.

After `Time` ms, `send_after/3` sends `Msg` to `Pid`.

If `Pid` is an atom, it is supposed to be the name of a registered process. The process referred to by the name is looked up at the time of delivery. No error is given if the name doesn't refer to a process. See also `start_timer/3` and `cancel_timer/1`.

Limitations: `Pid` must be a process on the local node. The timeout value must fit in 32 bits.

Failure: `badarg` if any arguments are of the wrong type, or do not obey the limitations noted above.

`erlang:set_cookie(Node, Cookie)`

Sets the "magic cookie" of `Node` to the atom `Cookie`. If `Node` is the current node, the BIF also sets the cookie of all other unknown nodes to `Cookie` (see `auth(3)`).

`setelement(Index, Tuple, Value)`

Returns a tuple which is a copy of the argument `Tuple` with the element given by the integer argument `Index` (the first element is the element with index 1) replaced by the argument `Value`.

```
> setelement(2, {10, green, bottles}, red).  
{10, red, bottles}
```

Failure: `badarg` if `Index` is not an integer, or `Tuple` is not a tuple, or if `Index` is less than 1 or greater than the size of `Tuple`.

`size(Item)`

Returns an integer which is the size of the argument `Item`, where `Item` must be either a tuple or a binary.

```
> size({morni, mulle, bwange}).  
3
```

Allowed in guard tests.

Failure: `badarg` if `Item` is not a tuple or a binary.

`spawn(Module, Function, ArgumentList)`

Returns the Pid of a new process started by the application of `Module:Function` to `ArgumentList`. *Note:* The new process created will be placed in the system scheduler queue and will be run some time later.

`error_handler:undefined_function(Module, Function, ArgumentList)` is evaluated by the new process if `Module:Function/Arity` does not exist (where `Arity` is the length of `ArgumentList`). The error handler can be redefined (see BIF `process_flag/2`). `Arity` is the length of the `ArgumentList`. If `error_handler` is undefined, or the user has redefined the default `error_handler` so its replacement is undefined, a failure with the reason `undef` will occur.

```
> spawn(speed, regulator, [high_speed, thin_cut]).
<0.13.1>
```

Failure: `badarg` if `Module` and/or `Function` is not an atom, or if `ArgumentList` is not a list.

`spawn(Node, Module, Function, ArgumentList)`

Works like `spawn/3`, with the exception that the process is spawned at `Node`. If `Node` does not exist, a useless Pid is returned.

Failure: `badarg` if `Node`, `Module`, or `Function` are not atoms, or `ArgumentList` is not a list.

`spawn_link(Module, Function, ArgumentList)`

This BIF is identical to the following code being evaluated in an atomic operation:

```
> Pid = spawn(Module, Function, ArgumentList),
  link(Pid),
  Pid.
```

This BIF is necessary since the process created might run immediately and fail before `link/1` is called.

Failure: See `spawn/3`.

`spawn_link(Node, Module, Function, ArgumentList)`

Works like `spawn_link/3`, except that the process is spawned at `Node`. If an attempt is made to spawn a process on a node which does not exist, a useless Pid is returned, and an `EXIT` signal will be received.

`spawn_opt(Module, Function, ArgumentList, Options)`

Works exactly like `spawn/3`, except that the following options can be given when creating the process:

`link` Sets a link to the parent process (`link spawn_link/3` does).

`{priority, Level}` Sets the priority of the new process. Equivalent to executing `process_flag(priority, Level)` in the start function of the new process, except that the priority will be set before the process is scheduled in the first time.

`{gc_switch, Threshold}` Sets the garbage collection strategy for the new process. If `Threshold` is 0, generational garbage collection will be used (this is default). If `Threshold` is greater than zero, the fullsweep algorithm will be used until the amount of live data (counted in words) reaches `Threshold`, at which point generational collection will be used. A large `Threshold` or the atom `infinity` means the fullsweep algorithm will be used.

Note: Generation collection is generally faster, especially for processes with a large amount of live data, but requires more memory since two heaps are allocated.

`{min_heap_size, Size}` Gives a minimum heap size in words. Setting this value higher than the system default might speed up some processes because less garbage collection is done. Setting too high value, however, might waste memory and slow down the system due to worse data locality. Therefore, it is recommended to use this option only for fine-tuning an application and to measure the execution time with various `Size` values.

`split_binary(Binary, Pos)`

Returns a tuple which contains two binaries which are the result of splitting `Binary` into two parts at position `Pos`. This is not a destructive operation. After this operation, there are three binaries altogether. Returns a tuple consisting of the two new binaries. For example:

```
1> B = list_to_binary("0123456789").
#Bin
2> size(B).
10
3> {B1, B2} = split_binary(B,3).
{#Bin, #Bin}
4> size(B1).
3
5> size(B2).
7
```

Failure: `badarg` if `Binary` is not a binary, or `Pos` is not an integer or is out of range.

`erlang:start_timer(Time, Pid, Msg)`

`Time` is a non-negative integer, `Pid` is either a pid or an atom, and `Msg` is any Erlang term. The function returns a reference.

After `Time` ms, `start_timer/3` sends the tuple `{timeout, Ref, Msg}` to `Pid`, where `Ref` is the reference returned by `start_timer/3`.

If `Pid` is an atom, it is supposed to be the name of a registered process. The process referred to by the name is looked up at the time of delivery. No error is given if the name doesn't refer to a process. See also `send_after/3` and `cancel_timer/1`.

Limitations: `Pid` must be a process on the local node. The timeout value must fit in 32 bits.

Failure: `badarg` if any arguments are of the wrong type, or do not obey the limitations noted above.

`statistics(Type)`

Returns information about the system. `Type` is an atom which is one of:

`run_queue` Returns the length of the run queue, that is the number of processes that are ready to run.

`runtime` Returns `{Total_Run_Time, Time_Since_Last_Call}`.

`wall_clock` Returns `{Total_Wallclock_Time, Wallclock_Time_Since_Last_Call}`.  
`wall_clock` can be used in the same manner as the atom runtime, except that real time is measured as opposed to runtime or CPU time.

`reductions` Returns `{Total_Reductions, Reductions_Since_Last_Call}`.

`garbage_collection` Returns `{Number_of_GC_s, Words_Reclaimed, 0}`. This information may not be valid for all implementations.

All times are in milliseconds.

```
> statistics(runtime).
{1690, 1620}
> statistics(reductions).
{2046, 11}
> statistics(garbage_collection).
{85, 23961, 0}
```

Failure: badarg if Type is not one of the atoms shown above.

`erlang:system_flag(Flag, Value)`

This BIF sets various system properties of the Erlang node. If `Flag` is a valid name of a system flag, its value is set to `Value`, and the old value is returned.

The currently allowed value for `Flag` is `keep_zombies`. The value of the `keep_zombies` flag is an integer which indicates how many processes to keep in memory when they exit, so that they can be inspected with `process_info`. Originally, the number is 0. Setting it to 0 disables the keeping of zombies. A negative number `-N` means to keep the `N` latest zombies; a positive value `N` means to keep the `N` first zombies. Setting the flag always clears away any already saved zombies. The maximum number of zombies which can be saved is 100. Resources owned by a zombie process are cleared away immediately when the process dies, for example ets tables and ports, and can not be inspected.

`term_to_binary(Term)`

This BIF returns the encoded value of any Erlang term and turns it into the Erlang external term format. It can be used for a variety of purposes, for example writing a term to a file in an efficient way, or sending an Erlang term to some type of communications channel not supported by distributed Erlang.

Returns a binary data object which corresponds to an external representation of the Erlang term `Term`.

`throw(Any)`

A non-local return from a function. If evaluated within a `catch`, `catch` will return the value `Any`.

```
> catch throw({hello, there}).
{hello, there}
```

Failure: `nocatch` if not evaluated within a `catch`.

time()

Returns the tuple {Hour, Minute, Second} of the current system time. The time zone correction is implementation-dependent.

```
> time().  
{9, 42, 44}
```

tl(List)

Returns List stripped of its first element.

```
> tl([geesties, guilies, beasties]).  
[guilies, beasties]
```

Failure: badarg if List is the empty list [], or is not a list. Allowed in guard tests.

trace(Pid, How, Flaglist)

Turns on tracing (if How == true) for the process Pid for all the trace items present in Flaglist. If How == false, the items in Flaglist are turned off. Pid must be a local process.

Flaglist can contain any number of the following atoms (the “message tags” refers to the list of message following below):

send Traces the messages the process Pid sends. Message tags: send, send\_to\_non\_existing\_process.

'receive' Traces the messages the process Pid receives. Message tags: 'receive'.

procs Traces process related events, for example spawn, link, exit. Message tags: spawn, exit, link, unlink, getting\_linked.

bifs Traces calls to BIFs. Message tags: call.

call Traces function calls. This flag only works for code that has been trace-compiled. Message tags: call, return.

running Traces scheduling of processes. Message tags: in, out.

garbage\_collection Traces garbage collections of processes. Message tags: gc\_start, gc\_end.

timestamp Make a time stamp in all trace messages. When a time stamp is used, traces of the form {Mod, Fun, Args} are converted to {Mod, Fun, Arity}, except when representing a BIF call. The time stamp (Ts) is of the same form as returned by erlang:now().

set\_on\_spawn Makes any process created by Pid inherit the flags of Pid, including the set\_on\_spawn flag.

set\_on\_first\_spawn Makes the first process created by Pid inherit the flags of Pid. That process does not inherit the set\_on\_first\_spawn flag.

set\_on\_link Makes any process linked by Pid inherit the flags of Pid, including the set\_on\_link flag.

set\_on\_first\_link Makes the first process linked to by Pid inherit the flags of Pid. That process does not inherit the set\_on\_first\_link flag.

suspend Suspend/resume the process.

The effect of combining `set_on_first_link` with `set_on_link` is the same as having `set_on_first_link` alone. Likewise for `set_on_spawn` and `set_on_first_spawn`.

The tracing process receives messages of the form:

```
{trace, 'receive', Pid, Message [,Ts]} When the traced Pid receives
something.
{trace, send, Msg, Pid, To [,Ts]} When Pid sends a message.
{trace, send_to_non_existing_process, Mes, Pid, To [,Ts]} When Pid sends a
message to a non existing process.
{trace, Pid, call, {M,F,A} [,Ts]} When Pid makes a function/BIF call. The
return values of calls are never supplied, only the call and its arguments.
{trace, Pid, return, {M,F,A}, Ts} When Pid returns to function {M,F,A}. This
trace message is only sent when both flag call and timestamp are used.
{trace, Pid, spawn, Pid2 [,Ts]} When Pid spawns a new process Pid2.
{trace, Pid, exit, Reason [,Ts]} When Pid exits with reason Reason.
{trace, Pid, link, Pid2 [,Ts]} When Pid links to a process Pid2.
{trace, Pid, unlink, Pid2 [,Ts]} When Pid removes the link from a process Pid2.
{trace, Pid, getting_linked, Pid2 [,Ts]} When Pid gets linked to a process
Pid2.
{trace, Pid, in, {M,F,A} [,Ts]} When Pid is scheduled to run. The process will
run in function {M,F,A}, where A is always the arity.
{trace, Pid, out, {M,F,A} [,Ts]} When Pid is scheduled out. The process was
running in function {M,F,A} where A is always the arity.
{trace, Pid, gc_start, Info, [,Ts]} Sent when garbage collection is about to be
started. Info is a list of two-element tuples, where the first element is a key, and
the second is the value. You should not depend on the tuples have any defined
order. Currently, the following keys are defined.
heap_size The size of the used part of the heap.
old_heap_size The size of the used part of the old heap.
stack_size The actual size of the stack.
recent_size The size of the data that survived the previous garbage collection.
mbuf_size The combined size of message buffers associated with the process.
All sizes are in words.
{trace, Pid, gc_end, Info, [,Ts]} Sent when garbage collection is finished. Info
contains the same kind of list as in the gc_start message, but the sizes reflect the
new sizes after garbage collection.
```

If the tracing process dies, the flags will be silently removed.

Only one process can trace a particular process. For this reason, attempts to trace an already traced process will fail.

Tracing of processes which trace other processes is not allowed.

Returns: `true`.

Failure: `badarg` if bad arguments are given.

`trunc(Number)`

Returns an integer by the truncation of `Number`. Allowed in guard tests.

```
> trunc(5.5).  
5
```

Failure: `badarg` if the argument is not a float, or an integer.

`tuple_to_list(Tuple)`

Returns a list which corresponds to `Tuple`. `Tuple` may contain any valid Erlang terms.

```
> tuple_to_list({share, {'Ericsson_B', 163}}).  
[share, {'Ericsson_B', 163}]
```

Failure: `badarg` if the argument is not a tuple.

`erlang:universaltime()`

Returns the current date and time according to Universal Time Coordinated (UTC), also called GMT, in the form `{{Year, Month, Day}, {Hour, Minute, Second}}` if supported by the underlying operating system. If not, `erlang:universaltime()` is equivalent to `erlang:localtime()`.

```
> erlang:universaltime().  
{{1996,11,6},{14,18,43}}
```

`erlang:universaltime_to_localtime(DateTime)`

Converts UTC date and time in `DateTime` to local date and time if supported by the underlying operating system. Otherwise, no conversion is done, and `DateTime` is returned. The return value is of the form `{{Year, Month, Day}, {Hour, Minute, Second}}`.

Failure: `badarg` if the argument is not a valid date and time tuple `{{Year, Month, Day}, {Hour, Minute, Second}}`.

```
> erlang:universaltime_to_localtime({{1996,11,6},{14,18,43}}).  
{{1996,11,7},{15,18,43}}
```

`unlink(Pid)`

Removes a link, if there is one, from the calling process to another process given by the argument `Pid`.

Returns `true`. Will not fail if not linked to `Pid`, or if `Pid` does not exist.

Failure: `badarg` if the argument is not a valid `Pid`.

`unregister(Name)`

Removes the registered name for a process, given by the atom argument `Name`.

Returns the atom `true`.

```
> unregister(db).  
true
```

Failure: `badarg` if `Name` is not the name of a registered process.

Users are advised not to unregister system processes.

`whereis(Name)`

Returns the Pid for the process registered under `Name` (see `register/2`). Returns undefined if no such process is registered.

```
> whereis(user).
```

```
<0.3.1>
```

Failure: `badarg` if the argument is not an atom.

# error\_handler (Module)

The error handler module defines what happens when certain types of errors occur.

## Exports

`undefined_function(Module, Func, ArgList) -> term()`

Types:

- `Module = Func = atom()`
- `ArgList = [term()]`

This function is evaluated if a call is made to `Module:Func(ArgList)` which is undefined. This function is evaluated inside the process making the original call.

If `Module` is interpreted, the interpreter is invoked and the return value of the interpreted `Func(ArgList)` call is returned.

Otherwise, it returns, if possible, the value of `apply(Module, Func, ArgList)` after an attempt has been made to autoload `Module`. If this is not possible, the function calling `Module:Func(ArgList)` is exited.

`undefined_lambda(Module, Fun, ArgList) -> term()`

Types:

- `Module = Fun = atom()`
- `ArgList = [term()]`

This function is evaluated if a call is made to `Fun(ArgList)` when the module defining the `fun` is not loaded. This function is evaluated inside the process making the original call.

If `Module` is interpreted, the interpreter is invoked and the return value of the interpreted `Fun(ArgList)` call is returned.

Otherwise, it returns, if possible, the value of `apply(Fun, ArgList)` after an attempt has been made to autoload `Module`. If this is not possible, the process calling the `fun` is exited.

## Notes

The code in `error_handler` is complex and should not be changed without fully understanding the interaction between the error handler, the `init` process of the code server, and the I/O mechanism of the code.

Changes in the code which may seem small can cause a deadlock as unforeseen consequences may occur. The use of `input` is dangerous in this type of code.

# error\_logger (Module)

The error logger is an event manager behaviour which runs with the registered name `error_logger` (see more about event managers/handlers in the Erlang Development Environment User's Guide and in `gen_event(3)`). All error messages from the Erlang runtime system are sent to this process as messages with the format `{emulator, Gleader, Str}`, where `Str` is a string which describes the error in plain English. The `Gleader` argument is the group leader process of the process causing the error. This is useful in a distributed setting as all error messages can be returned to the `error_logger` process on the originating node.

All errors detected by the standard libraries are reported with the `error_logger` functions. Errors detected in application modules should also be reported through the `error_logger` in order to get uniform reports.

Associated event handlers can be used to add private `types` of reports to the `error_logger`. An event handler which recognizes the specialized report type is first added to the `error_logger` (`add_report_handler/1,2`)

The standard configuration of the `error_logger` supports the logging of errors to the `tty`, or to a specified `file`. There is also a multi-file logger which logs all events, not only the standard error events, to several files. (see `log_mf_h(3)`).

All error events are tagged with the group leader `Gleader` in order to send the error to the originating node.

## Exports

```
start() -> {ok, Pid} | {error, What}
```

```
start_link() -> {ok, Pid} | {error, What}
```

Types:

- `Pid = pid()`
- `What = {already_started, Pid} | term()`

Starts the `error_logger`. The `start_link` function should be used when the `error_logger` is supervised

```
error_report(Report) -> ok
```

Types:

- `Report = [{Tag, Data}] | [term()] | string() | term()`
- `Tag = term()`
- `Data = term()`

Sends a standard error report event to the error logger. This report event is handled by the standard event handler. The report is formatted as follows:

```
Tag1:   Data1
Tag2:   Data2
Term1
Term2
```

If Report is a string(), the string is written.

The report is written with an error heading.

```
error_report(Type,Report) -> ok
```

Types:

- Type = term()
- Report = [{Tag, Data}] | [term()] | string() | term()
- Tag = term()
- Data = term()

Sends a user defined error report type event to the error logger. If specialized error handling is required, an event handler recognizing this Type of report must first be added to the error\_logger.

It is recommended that the Report follows the same structure as error\_report/1 above.

```
info_report(Report) -> ok
```

Types:

- Report = [{Tag, Data}] | [term()] | string() | term()
- Tag = term()
- Data = term()

Sends an information report to the error logger. This report event is handled by the standard event handler. The report is formatted as follows:

```
Tag1:   Data1
Tag2:   Data2
Term1
Term2
```

If Report is a string(), the string is written.

The report is written with an information heading.

```
info_report(Type,Report) -> ok
```

Types:

- Type = term()
- Report = [{Tag, Data}] | [term()] | string() | term()
- Tag = term()
- Data = term()

Sends a user defined information report type event to the error logger. If specialized error handling is required, an event handler recognizing this Type of report must first be added to the `error_logger`.

It is recommended that the Report follows the same structure as `info_report/1` above.

`error_msg(Format) -> ok`

`error_msg(Format,Args) -> ok`

`format(Format,Args) -> ok`

Types:

- `Format = string()`
- `Args = [term()]`

Sends an error event to the error logger. The `Format` and `Args` arguments are the same as the arguments of `io:format/2`. These events are handled by the standard event handler.

`info_msg(Format) -> ok`

`info_msg(Format,Args) -> ok`

Types:

- `Format = string()`
- `Args = [term()]`

Sends an information event to the error logger. The `Format` and `Args` arguments are the same as the arguments of `io:format/2`. These events are handled by the standard event handler.

`tty(Flag) -> ok`

Types:

- `Flag = true | false`

Enables or disables error printouts to the tty. If `Flag` is `false`, all text that the error logger would have sent to the terminal is discarded. If `Flag` is `true`, error messages are sent to the terminal screen.

`logfile(Request) -> ok | FileName | {error, What}`

Types:

- `Request = {open, FileName} | close | filename`
- `FileName = atom() | string()`
- `What = term()`

This function makes it possible to append a copy of all standard error printouts to a file. It can be used in combination with the `tty(false)` function in to have a silent system, where all errors are logged to a file.

Request can be:

- `{open, Filename}`. Opens the file `Filename` to store a copy of all error messages. Returns `ok`, or `{error, What}`.
- `close`. Closes the current log file. Returns `ok`, or `{error, What}`.

- `filename`. Returns `{error, What}` or `FileName`, where `FileName` is the name of the open log file.

There can only be one active log file.

`add_report_handler(Module) -> ok | Other`

`add_report_handler(Module,Args) -> ok | Other`

Types:

- `Module = atom()`
- `Args = term()`
- `Other = term()`

Adds a new event handler to the error logger. The event handler is initialized by a call to the `Module:init/1` function. This function must return `{ok, State}`. If anything else (`Other`) is returned, the handler is not added.

The report (event) handler will be called for every error event that the error logger receives (`Module:handle_event/2`). Errors dedicated to this handler should be handled accordingly.

`delete_report_handler(Module) -> Return | {error, What}`

Types:

- `Module = atom()`
- `Return = term()`
- `What = term()`

Deletes an error report (event) handler. The `Module:terminate/2` function is called in order to finalize the event handler. The return value of the `terminate/2` function is `Return`.

`swap_handler(ToHandler) -> ok`

Types:

- `ToHandler = tty | {logfile, File}`
- `File = atom() | string()`

The `error_logger` event manager is initially started with a primitive event handler which buffers and prints the raw error events. However, this function does install the standard event handler to be used according to the system configuration.

## Events

The error logger event manager forwards the following events to all added event handlers. In the events that follow, `Gleader` is the group leader process identity of the error reporting process, and `EPid` is the process identity of the `error_logger`. All other variables are described with the function in which they appear.

{`error_report`, `Gleader`, {`EPid`, `std_error`, `Report`}} This event is generated when the `error_report/1` function is called.

{`error_report`, `Gleader`, {`EPid`, `Type`, `Report`}} This event is generated when the `error_report/2` function is called.

{`info_report`, `Gleader`, {`EPid`, `std_info`, `Report`}} This event is generated when the `info_report/1` function is called.

{`info_report`, `Gleader`, {`EPid`, `Type`, `Report`}} This event is generated when the `info_report/2` function is called.

{`error`, `Gleader`, {`EPid`, `Format`, `Args`}} This event is generated when the `error_msg` or `format` functions are called.

{`info_msg`, `Gleader`, {`EPid`, `Format`, `Args`}} This event is generated when the `info_msg` functions are called.

{`info`, `Gleader`, {`EPid`, `term()`, []}} This structure is only used by the `init` process for erroneously received messages.

{`emulator`, `Gleader`, `string()`} This event is generated by the runtime system. If the error was not issued by a special process, `Gleader` is `noproc`. This event should be handled in the `handle_info/2` function of the event handler.

### Note:

All events issued by a process which has the group leader `Gleader` process located on another node will be passed to this node by the `error_logger`.

## See Also

`gen_event(3)`, `log_mf_h(3)`

# file (Module)

The module `file` provides an interface to the file system.

Most functions have a name argument such as a file name or directory name, which is either an atom, a string, or a deep list of characters and atoms. A path is a list of directory names. If the functions are successful, they return `ok`, or `{ok, Value}`.

If an error occurs, the return value has the format `{error, Reason}`. Reason is an atom which is named from the Posix error codes used in Unix, and in the runtime libraries of most C compilers. In the following descriptions of functions, the most typical error codes are listed. By matching the error code, applications can use this information for error recovery. To produce a readable error string, use `format_error/1`.

## Exports

`read_file(Filename)`

Returns `{ok, Binary}`, where Binary is a binary data object that contains the contents of `Filename`, or `{error, Reason}` if an error occurs.

Typical error reasons:

`enoent` The file does not exist.

`eaccess` Missing permission for reading the file, or for searching one of the parent directories.

`eisdir` The named file is a directory.

`enotdir` A component of the file name is not a directory. On some platforms, `enoent` is returned instead.

`enomem` There is not enough memory for the contents of the file.

`write_file(Filename, Binary)`

Writes the contents of the binary data object `Binary` to the file `Filename`. The file is created if it does not exist already. If it exists, the previous contents are overwritten. Returns `ok`, or `{error, Reason}`.

Typical error reasons are:

`enoent` A component of the file name does not exist.

`enotdir` A component of the file name is not a directory. On some platforms, `enoent` is returned instead.

`enospc` There is a no space left on the device.

`eaccess` Missing permission for writing the file or searching one of the parent directories.

`eisdir` The named file is a directory.

`get_cwd()`

Returns `{ok, CurDir}`, where `CurDir` (a string) is the current working directory of the file server.

**Note:**

In rare circumstances, this function can fail on Unix. It may happen if read permission does not exist for the parent directories of the current directory.

Typical error reasons are:

`eaccess` Missing read permission for one of the parents of the current directory.

`get_cwd(Drive)`

`Drive` should be of the form "Letter:", for example "c:". Returns `{ok, CurDir}` or `{error, Reason}`, where `CurDir` (a string) is the current working directory of the drive specified.

This function returns `{error, enotsup}` on platforms which have no concept of current drive (Unix, for example).

Typical error reasons are:

`enotsup` The operating system have no concept of drives.

`eaccess` The drive does not exist.

`EINVAL` The format of `Drive` is invalid.

`set_cwd(DirName)`

Sets the current working directory of the file server to `DirName`. Returns `ok` if successful.

Typical error reasons are:

`ENOENT` The directory does not exist.

`ENOTDIR` A component of `DirName` is not a directory. On some platforms, `ENOENT` is returned.

`EACCESS` Missing permission for the directory or one of its parents.

`delete(Filename)`

Tries to delete the file `Filename`. Returns `ok` if successful.

Typical error reasons are:

`ENOENT` The file does not exist.

`EACCESS` Missing permission for the file or one of its parents.

`EPERM` The file is a directory and the user is not super-user.

`enotdir` A component of the file name is not a directory. On some platforms, `enoent` is returned instead.

`rename(Source, Destination)`

Tries to rename the file `Source` to `Destination`. It can be used to move files (and directories) between directories, but it is not sufficient to specify the destination only. The destination file name must also be specified. For example, if `bar` is a normal file and `foo` and `baz` are directories, `rename("foo/bar", "baz")` returns an error, but `rename("foo/bar", "baz/bar")` succeeds. Returns `ok` if it is successful.

**Note:**

Renaming of open files is not allowed on most platforms (see `eaccess` below).

Typical error reasons:

`eaccess` Missing read or write permissions for the parent directories of `Source` or `Destination`. On some platforms, this error is given if either `Source` or `Destination` is open.

`eexist` `Destination` is not an empty directory.

`EINVAL` `Source` is a root directory, or `Destination` is a sub-directory of `Source`.

`EISDIR` `Destination` is a directory, but `Source` is not.

`ENOENT` `Source` does not exist.

`ENOTDIR` `Source` is a directory, but `Destination` is not.

`EXDEV` `Source` and `Destination` are on different file systems.

`make_dir(DirName)`

Tries to create the directory `DirName`. Missing parent directories are NOT created. Returns `ok` if successful.

Typical error reasons are:

`eaccess` Missing search or write permissions for the parent directories of `DirName`.

`eexist` There is already a file or directory named `DirName`.

`ENOENT` A component of `DirName` does not exist.

`ENOSPC` There is a no space left on the device.

`ENOTDIR` A component of `DirName` is not a directory. On some platforms, `ENOENT` is returned instead.

`del_dir(DirName)`

Tries to delete the directory `DirName`. The directory must be empty before it can be deleted. Returns `ok` if successful.

Typical error reasons are:

`eaccess` Missing search or write permissions for the parent directories of `DirName`.

`eexist` The directory is not empty.  
`enoent` The directory does not exist.  
`enotdir` A component of `DirName` is not a directory. On some platforms, `enoent` is returned instead.  
`EINVAL` Attempt to delete the current directory. On some platforms, `EACCES` is returned instead.

`make_link(Existing, New)`

Makes a hard link from `Existing` to `New`, on platforms that support links (Unix). This function returns `ok` if the link was successfully created, or `{error, Reason}`. On platforms that don't support links, `{error, enotsup}` will be returned.

Typical error reasons:

`EACCES` Missing read or write permissions for the parent directories of `Existing` or `New`.  
`EEXIST` `new` already exists.  
`ENOTSUP` Hard links are not supported on this platform.

`make_symlink(Name1, Name2)`

This function creates a symbolic link `Name2` to the file or directory `Name1`, on platforms that support symbolic links (most Unix systems). `Name1` need not exist. This function returns `ok` if the link was successfully created, or `{error, Reason}`. On platforms that don't support symbolic links, `{error, enotsup}` will be returned.

Typical error reasons:

`EACCES` Missing read or write permissions for the parent directories of `Existing` or `New`.  
`EEXIST` `new` already exists.  
`ENOTSUP` Symbolic links are not supported on this platform.

`list_dir(DirName)`

Lists all the files in a directory. Returns `{ok, FilenameList}` if successful. Otherwise, it returns `{error, Reason}`. `FilenameList` is a list of the names of all the files in the directory. Each name is a string. The names are not sorted.

Typical error reasons are:

`EACCES` Missing search or write permissions for `DirName` or one of its parent directories.  
`ENOENT` The directory does not exist.

`read_file_info(Filename)`

Retrieves information about a file. Returns `{ok, FileInfo}` if successful, otherwise `{error, Reason}`. `FileInfo` is a record. Its definition can be found by including `file.hrl` from the kernel application:

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

The record contains the following fields.

- size** Size of file in bytes.
- type** The type of the file which can be device, directory, regular, or other.
- access** The current system access to the file, which is one of the atoms read, write, read\_write, or none.
- atime** The last (local) time the file was read, in the format {{Year, Month, Day}, {Hour, Minute, Second}}.
- mtime** The last (local) time the file was written, in the format {{Year, Month, Day}, {Hour, Minute, Second}}.
- ctime** The interpretation of this time field depends on the operating system. On Unix, it is the last time the file or the inode was changed. In Windows, it is the create time. The format is {{Year, Month, Day}, {Hour, Minute, Second}}.
- mode** An integer which gives the file permissions as a sum of the following bit values:
- 8#00400** read permission: owner
  - 8#00200** write permission: owner
  - 8#00100** execute permission: owner
  - 8#00040** read permission: group
  - 8#00020** write permission: group
  - 8#00010** execute permission: group
  - 8#00004** read permission: other
  - 8#00002** write permission: other
  - 8#00001** execute permission: other
  - 16#800** set user id on execution
  - 16#400** set group id on execution
- On Unix platforms, other bits than those listed above may be set.
- links** Number of links to the file (this will always be 1 for file systems which have no concept of links).
- major\_device** An integer which identifies the file system where the file is located. In Windows, the number indicates a drive as follows: 0 means A:, 1 means B:, and so on.
- minor\_device** Only valid for character devices on Unix. In all other cases, this field is zero.
- inode** An integer which gives the inode number. On non-Unix file systems, this field will be zero.
- uid** An integer which indicates the owner of the file. Will be zero for non-Unix file systems.
- gid** An integer which gives the group that the owner of the file belongs to. Will be zero for non-Unix file systems.

Typical error reasons:

- eaccess** Missing search permission for one of the parent directories of the file.
- enoent** The file does not exist.
- enotdir** A component of the file name is not a directory. On some platforms, **enoent** is returned instead.

`read_link_info(Filename)`

This function works like `read_file_info/1`, except that if `Filename` is a symbolic link, information about the link will be returned in the `file_info` record and the `type` field of the record will be set to `symlink`. If `Filename` is not a symbolic link, this function returns exactly the same result as `read_file_info/1`. On platforms that don't support symbolic link, this function is always equivalent to `read_file_info/1`.

`read_link(Linkname)`

This function returns `{ok,Filename}` if `Linkname` refers to a symbolic link or `{error,Reason}` otherwise. On platforms that don't support symbolic links, the return value will be `{error,enotsup}`.

Typical error reasons:

`EINVAL` `Linkname` does not refer to a symbolic link.

`ENOENT` The file does not exist.

`ENOTSUP` Symbolic links are not supported on this platform.

`file_info(Filename)`

**Note:**

This function is obsolete. Use `read_file_info` instead.

Retrieves information about a file. Returns `{ok, FileInfo}` if successful, otherwise `{error, Reason}`. `FileInfo` is a tuple with the following fields:

`{Size,Type,Access,AccessTime,ModifyTime,UnUsed1,UnUsed2}`

`Size` The size of the file in bytes.

`Type` The type of file which is device, directory, regular, or other.

`Access` The current system access to the file, which is one of the atoms `read`, `write`, `read_write`, or `none`.

`AccessTime` The last time the file was read, shown in the format `{Year, Month, Day, Hour, Minute, Second}`.

`ModifyTime` The last time the file was written, shown in the format `{Year, Month, Day, Hour, Minute, Second}`.

`UnUsed1`, `UnUsed2` These fields are not used, but reserved for future expansion. They probably contain unused.

Typical error reasons: Same as for `read_file_info/1`.

`write_file_info(Filename, FileInfo)`

Change file information. Returns `ok` if successful, otherwise `{error, Reason}`.

`FileInfo` is a record. Its definition can be found by including `file.hrl` from the kernel application:

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

The following fields are used from the record if they are given.

**atime** The last (local) time the file was read, in the format `{{Year, Month, Day}, {Hour, Minute, Second}}`.

**mtime** The last (local) time the file was written, in the format `{{Year, Month, Day}, {Hour, Minute, Second}}`.

**ctime** On Unix, any value give for this field will be ignored (the “ctime” for the file will be set to the current time). On Windows, this field is the new creation time to set for the file. The format is `{{Year, Month, Day}, {Hour, Minute, Second}}`.

**mode** An integer which gives the file permissions as a sum of the following bit values:

- 8#00400** read permission: owner
- 8#00200** write permission: owner
- 8#00100** execute permission: owner
- 8#00040** read permission: group
- 8#00020** write permission: group
- 8#00010** execute permission: group
- 8#00004** read permission: other
- 8#00002** write permission: other
- 8#00001** execute permission: other
- 16#800** set user id on execution
- 16#400** set group id on execution

On Unix platforms, other bits than those listed above may be set.

**uid** An integer which indicates the owner of the file. Ignored for non-Unix file systems.

**gid** An integer which gives the group that the owner of the file belongs to. Ignored non-Unix file systems.

Typical error reasons:

**eaccess** Missing search permission for one of the parent directories of the file.

**enoent** The file does not exist.

**enotdir** A component of the file name is not a directory. On some platforms, **enoent** is returned instead.

`change_owner(Filename, Uid)`

Change owner of a file. See `write_file_info/2`.

`change_owner(Filename, Uid, Gid)`

Change owner and group of a file. See `write_file_info/2`.

`change_group(Filename, Gid)`

Change group of a file. See `write_file_info/2`.

`change_time(Filename, Mtime)`

Change the modification and access times of a file. See `write_file_info/2`.

`change_time(Filename, Mtime, Atime)`

Change the modification and access times of a file. See `write_file_info/2`.

`consult(Filename)`

Opens file `Filename` and reads all the Erlang terms in it. Returns one of the following:

- `{ok, TermList}` The file was successfully read.
- `{error, Atom}` An error occurred when opening the file or reading it. The `Atom` is a Posix error code. See the description of `open/2` for a list of typical error codes.
- `{error, {Line, Mod, Term}}` An error occurred when interpreting the Erlang terms in the file. Use the `format_error/1` function to convert the three-element tuple to an English description of the error.

`path_consult(Path, Filename)`

Searches the path `Path` (a list of directory names) until the file `Filename` is found. If `Filename` is an absolute file name, `Path` is ignored. The file is opened and all the terms in it are read. The function returns one of the following:

- `{ok, TermList, FullName}` The file was successfully read. `FullName` is the full name of the file which was opened and read.
- `{error, enoent}` The file could not be found in any of the directories in `Path`.
- `{error, Atom}` An error occurred when opening the file or reading it. The `Atom` is a Posix error code. See the description of `open/2` for a list of typical error codes.
- `{error, {Line, Mod, Term}}` An error occurred when interpreting the Erlang terms in the file. Use the `format_error/1` function to convert the three-element tuple to an English description of the error.

`eval(Filename)`

Opens the file `Filename` and evaluates all the expression sequences in it. It returns one of the following:

- `ok` The file was read and evaluated. The actual result of the evaluation is not returned; any expression sequence in the file must be there for its side effect.
- `{error, Atom}` An error occurred when opening the file or reading it. The `Atom` is a Posix error code. See the description of `open/2` for a list of typical error codes.
- `{error, {Line, Mod, Term}}` An error occurred when interpreting the Erlang terms in the file. Use the `format_error/1` function to convert the three-element tuple to an English description of the error.

`path_eval(Path, Filename)`

Searches the path `Path` (a list of directory names) until the file `Filename` is found. If `Filename` is an absolute file name, `Path` is ignored. The file is opened and all the expression sequences in it are evaluated. The function returns one of the following:

- `{ok, FullName}` The file was read. `FullName` is the full name of the file which was opened and evaluated.
- `{error, enoent}` The file could not be found in any of the directories in `Path`.

- `{error, Atom}` An error occurred when opening the file or reading it. The `Atom` is a Posix error code. See the description of `open/2` for a list of typical error codes.
- `{error, {Line, Mod, Term}}` An error occurred when interpreting the Erlang terms in the file. Use the `format_error/1` function to convert the three-element tuple to an English description of the error.

`open(Filename, ModeList)`

Opens the file `Filename` in the mode determined by `ModeList`. `ModeList` may contain one or more of the following items:

- `read` The file, which must exist, is opened for reading.
- `write` The file is opened for writing. It is created if it does not exist. Otherwise, it is truncated (unless combined with `read`).
- `append` The file will be opened for writing, and it will be created if it does not exist. Every write operation to a file opened with `append` will take place at the end of the file.
- `raw` The `raw` option allows faster access to a file, because no Erlang process is needed to handle the file. However, a file opened in this way has the following limitations:
- The functions in the `io` module cannot be used, because they can only talk to an Erlang process. Instead, use the `read/2` and `write/2` functions.
  - Only the Erlang process which opened the file can use it.
  - A remote Erlang file server cannot be used; the computer on which the Erlang node is running must have access to the file system (directly or through NFS).
- `binary` This option can only be used if the `raw` option is specified as well. When specified, read operations on the file using the `read/2` function will return binaries rather than lists.

If both `read` and `write` are specified, the file is created if it does not exist. It is not truncated if it exists.

Returns:

- `{ok, IoDevice}` The file has been opened in the requested mode. `IoDevice` is a reference to the file.
- `{error, Reason}` The file could not be opened.

A file descriptor is the `Pid` of the process which handles the file. The file process is linked to the process which originally opened the file. If any process to which the file process is linked terminates, the file will be closed by the file process and the process itself will be terminated. The file descriptor returned from this call can be used as an argument to the I/O functions (see `io`).

**Note:**

In previous versions of `file`, modes were given as one of the atoms `read`, `write`, or `read_write` instead of a list. This is still allowed for reasons of backwards compatibility, but should not be used for new code. Also note that `read_write` is not allowed in a mode list.

Typical error reasons:

`enoent` The file does not exist.  
`eaccess` Missing permission for reading the file or searching one of the parent directories.  
`eisdir` The named file is a directory.  
`enotdir` A component of the file name is not a directory. On some platforms, `enoent` is returned instead.  
`enospc` There is a no space left on the device (if `write` access was specified).

`path_open(Path, Filename, Mode)`

Searches the path `Path` (a list of directory names) until the file `Filename` is found. If `Filename` is an absolute file name, `Path` is ignored. The function returns one of the following:

`{ok, IoDevice, FullName}` The file was opened in the requested mode. `IoDevice` is a reference to the file and `FullName` is the full name of the file which was opened.  
`{error, enoent}` `Filename` was not found in the path.  
`{error, Reason}` There was an error opening `Filename`.

`close(IoDevice)`

Closes the file referenced by `IoDevice`. It returns `ok`.

`read(IoDevice, Number)`

Reads `Number` bytes from the file described by `IoDevice`. This function is the only way to read from a file opened in raw mode (although it works for normally opened files, too). Returns:

`{ok, ListOrBinary}` If the file was opened in binary mode, the read bytes are returned in a binary, otherwise in a list. The list or binary will be shorter than the the number of bytes requested if the end of the file is reached.  
`eof` `eof` is returned if the `Number` was greater than zero and end of file was reached before anything at all could be read.  
`{error, Reason}` A Posix error code will be returned if an error occurred.  
Typical error reasons:  
`ebadf` The file is not opened for reading.

`write(IoDevice, Bytes)`

Writes `Bytes` (possibly a deep list of characters, or a binary) to the file described by `IoDevice`. This function is the only way to write to a file opened in raw mode (although it works for normally opened files, too).

This function returns `ok` if successful, and `{error, Reason}` otherwise.

Typical error reasons are:

`ebadf` The file is not opened for writing.  
`enospc` There is a no space left on the device.

`position(IoDevice, Location)`

Sets the position of the file referenced by `IoDevice` to `Location`. Returns `{ok, NewPosition}` (as absolute offset) if successful, otherwise `{error, Reason}`. `Location` is one of the following:

`{bof, Offset}` Absolute offset

`{cur, Offset}` Offset from the current position

`{eof, Offset}` Offset from the end of file

`Integer` The same as `{bof, Integer}`

`bof || cur || eof` The same as above with `Offset 0`.

Typical error reasons are:

`EINVAL` Either the `Location` was illegal, or it evaluated to a negative offset in the file.

Note that if the resulting position is a negative value you will get an error but after the call it is undefined where the file position will be.

`pread(IoDevice, Location, Number)`

Combines `position/2` and `read/2` in one operation, which is more efficient than calling them one at a time. If `IoDevice` has been opened in raw mode, some restrictions apply: `Location` is only allowed to be an integer; and the current position of the file is undefined after the operation.

`pwrite(IoDevice, Location, Bytes)`

Combines `position/2` and `write/2` in one operation, which is more efficient than calling them one at a time. If `IoDevice` has been opened in raw mode, some restrictions apply: `Location` is only allowed to be an integer; and the current position of the file is undefined after the operation.

`truncate(IoDevice)`

Truncates the file referenced by `IoDevice` at the current position. Returns `ok` if successful, otherwise `{error, Reason}`.

`sync(IoDevice)`

Makes sure that any buffers kept by the operating system (not by the Erlang system) are written to disk. On some platforms, this function might have no effect.

`format_error(ErrorDescriptor)`

Given the error reason returned by any function in this module, it returns a descriptive string of the error in English.

## POSIX Error Codes

**eaccess** permission denied  
**again** resource temporarily unavailable  
**badf** bad file number  
**ebusy** file busy  
**edquot** disk quota exceeded  
**eexist** file already exists  
**efault** bad address in system call argument  
**efbig** file too large  
**eintr** interrupted system call  
**EINVAL** invalid argument  
**EIO** I/O error  
**EISDIR** illegal operation on a directory  
**ELOOP** too many levels of symbolic links  
**EMFILE** too many open files  
**ENLINK** too many links  
**ENAMETOOLONG** file name too long  
**ENFILE** file table overflow  
**ENODEV** no such device  
**ENOENT** no such file or directory  
**ENOMEM** not enough memory  
**ENOSPC** no space left on device  
**ENOTBLK** block device required  
**ENOTDIR** not a directory  
**ENOTSUP** operation not supported  
**ENXIO** no such device or address  
**EPERM** not owner  
**EPIPE** broken pipe  
**EROFS** read-only file system  
**ESPIPE** invalid seek  
**ESRCH** no such process  
**ESTALE** stale remote file handle  
**EXDEV** cross-domain link

## Warnings

If an error occurs when accessing an open file with the `io` module, the process which handles the file will exit. The dead file process might hang if a process tries to access it later. This will be fixed in a future release.

## See Also

filename(3)

## gen\_tcp (Module)

The `gen_tcp` module provides functions for communicating with sockets using the TCP/IP protocol.

The available options are described in the `setopts/2` [page 125] function in the `inet` manual page.

The possible `{error, Reason}` results are described in the `inet` [page 126] manual page.

The following code fragment provides a simple example of a client connecting to a server at port 5678, transferring a binary and closing the connection.

```
client() ->
  SomeHostInNet = "localhost" % to make it runnable on one machine
  {ok, Sock} = gen_tcp:connect(SomeHostInNet, 5678,
                              [binary, {packet, 0}]),
  ok = gen_tcp:send(Sock, "Some Data"),
  ok = gen_tcp:close(Sock).
```

At the other end a server is listening on port 5678, accepts the connection and receives the binary.

```
server() ->
  {ok, LSock} = gen_tcp:listen(5678, [binary, {packet, 0},
                                     {active, false}]),
  {ok, Sock} = gen_tcp:accept(LSock),
  {ok, Bin} = do_recv(Sock, []),
  ok = gen_tcp:close(Sock),
  Bin.

do_recv(Sock, Bs) ->
  case gen_tcp:recv(Sock, 0) of
    {ok, B} ->
      do_recv(Sock, [Bs, B]);
    {error, closed} ->
      {ok, list_to_binary(Bs)}
  end.
```

## Exports

`accept(ListenSocket) -> {ok, Socket} | {error, Reason}`

`accept(ListenSocket, Timeout) -> {ok, Socket} | {error, Reason}`

Types:

- ListenSocket = socket()
- Socket = socket()
- Timeout = integer()
- Reason = atom()

Accepts an incoming connection request on a listen socket. Socket must be a socket returned from `listen/1`. If no Timeout argument is specified, or it is `infinity`, the `accept` function will not return until a connection has been established. If Timeout is specified and no connection is accepted within the given time, `accept` will return `{error, timeout}`.

Packets can be sent to the returned socket using the `send/2` function. Packets sent from the peer will be delivered as messages

`{tcp, Socket, Data}`

unless `{active, false}` was specified in the option list for the listen socket, in which case packets should be retrieved by calling `recv/2`.

`close(Socket) -> ok | {error, Reason}`

Types:

- Socket = socket()
- Reason = atom()

Closes an TCP socket.

`connect(Address, Port, Options) -> {ok, Socket} | {error, Reason}`

`connect(Address, Port, Options, Timeout) -> {ok, Socket} | {error, Reason}`

Types:

- Address = string() | atom() | ip\_address()
- Port = Timeout = integer()
- Options = list()
- Socket = socket()
- Reason = atom()

Connects to a server on TCP port Port on the host with IP address Address. The Address argument can be either a hostname, or an IP address.

The available options are:

`list` Received Packet is delivered as a list.

`binary` Received Packet is delivered as a binary.

**common inet options** The common inet options available are described in the `setopts/2` [page 125] function in the `inet` manual page.

Packets can be sent to the returned socket using the `send/2` function. Packets sent from the peer will be delivered as messages

```
{tcp, Socket, Data}
```

If the socket was closed the following message is delivered:

```
{tcp_closed, Socket}
```

If an error occurred on the socket the following message is delivered:

```
{tcp_error, Socket, Reason}
```

unless the socket is in passive mode, in which case packets are retrieved by calling `recv/2`.

The optional `Timeout` parameter specifies a timeout in milliseconds. The default value is infinity.

```
controlling_process(Socket, NewOwner) -> ok | {error, eperm}
```

Types:

- `Socket` = `socket()`
- `NewOwner` = `pid()`

Assigns a new controlling process to `Socket`. The controlling process is the process which will receive messages from the socket. If called by any other process than the current owner `{error, eperm}` will be returned.

```
listen(Port, Options) -> {ok, Socket} | {error, Reason}
```

Types:

- `Port` = `integer()`
- `Options` = `list()`
- `Socket` = `socket()`
- `Reason` = `atom()`

Sets up socket to listen on the port `Port` on the local host.

If the port number is zero, the `listen` function picks an available port number (use `inet:port/1` to retrieve it); otherwise, the specified port number is used.

The available options are described in the `setopts/2` [page 125] function in the `inet` manual page. Additionally, the option `{backlog, B}` can be given, where `B` is an integer  $\geq 0$ . The backlog value defaults to 5. The backlog value defines the maximum length the queue of pending connections may grow to.

The returned socket can only be used in calls to `accept`.

```
recv(Socket, Length) -> {ok, Packet} | {error, Reason}
```

```
recv(Socket, Length, Timeout)
```

Types:

- `Socket` = `socket()`
- `Length` = `integer()`
- `Packet` = `list() | binary()`

- Timeout = integer()
- Reason = atom()

This function receives a packet from a socket in passive mode. A closed socket is indicated by a return value of {error, closed}.

The Length argument is only meaningful when the socket is in raw mode and denotes number of bytes to read. If Length = 0 all available bytes are returned.

The optional Timeout parameter specifies a timeout in milliseconds. The default value is infinity.

send(Socket, Packet) -> ok | {error, Reason}

Types:

- Socket = socket()
- Packet = list() | binary()
- Reason = atom()

Sends a packet on a socket.

# gen\_udp (Module)

The `gen_udp` module is an interface to User Datagram Protocol (UDP).

The possible `{error, Reason}` results are described in the `inet` [page 126] manual page.

## Exports

`close(Socket) -> ok | {error, Reason}`

Types:

- `Socket = Reason = term()`

Removes the `Socket` created with `open/1` or `open/2`.

`controlling_process(Socket, NewOwner) ->`

Types:

- `Socket = term()`
- `NewOwner = pid()`

The process associated with a `Socket` is changed to `NewOwner`. The `NewOwner` will receive all subsequent data.

`open(Port) -> {ok, Socket} | {error, Reason}`

`open(Port, Options) -> {ok, Socket} | {error, Reason}`

Types:

- `Port = integer(0..65535)`
- `Options = list()`
- `Socket = term()`
- `Reason = term()`

Associates a UDP port number (`Port`) with the calling process. It returns `{ok, Socket}`, or `{error, Reason}`. The returned `Socket` is used to send packets from this port with the `send/4` function. `Options` is a list of options associated with this port.

When UDP packets arrive at the opened `Port` they will be delivered as messages of the type `{udp, Socket, IP, InPortNo, Packet}`

`IP` and `InPortNo` define the address from which `Packet` came. `Packet` is a list of bytes if the option `list` was specified. `Packet` is a binary if the option `binary` was specified.

The available options are:

`list` Received `Packet` is delivered as a list.

binary Received Packet is delivered as a binary.

**common inet options** The common inet options available are described in the `setopts/2` [page 125] function in the `inet` manual page.

**Note:**

The common inet option `{active, false}` which sets passive mode is not available in this version, but will be implemented in a following version.

If you set `Port` to 0, the underlying Operating System assigns a free UDP port. (You can find out which port by calling `inet:port(Socket)`.)

If any of the following functions are called with a `Socket` that was not opened by the calling process, they will return `{error, not_owner}`. The ownership of a `Socket` can be transferred to another process with `controlling_process/2`.

`recv(Socket, Length) -> {ok, Packet} | {error, Reason}`

`recv(Socket, Length, Timeout)`

Types:

- `Socket = socket()`
- `Length = integer()`
- `Packet = list() | binary()`
- `Timeout = integer()`
- `Reason = atom()`

This function receives a packet from a socket in passive mode.

The optional `Timeout` parameter specifies a timeout in milliseconds. The default value is `infinity`.

`send(S, Address, Port, Packet) -> ok | {error, Reason}`

Types:

- `Address = { integer(), integer(), integer(), integer() } | atom() | string()`
- `Port = integer(0..65535)`
- `Packet = [byte()] | binary()`
- `Reason = term()`

Sends `Packet` to the specified address (`address, port`). It returns `ok`, or `{error, Reason}`. `Address` can be an IP address expressed as a tuple, for example `{192, 0, 0, 1}`. It can also be a host name expressed as an atom or a string, for example `'somehost.some.domain'`. `Port` is an integer, and `Packet` is either a list of bytes or a binary.

# global (Module)

This documentation describes the Global module which consists of the following functionalities:

1. Registration of global names
2. Global locks
3. Monitoring nodes
4. Maintenance of the fully connected network

These services are controlled via the process `global` which exists on every node. `global` is started automatically when a node is started.

The ability to globally register names is a central concept in the programming of distributed Erlang systems. In this module, the equivalent of the `register/2` and `whereis/1` BIFs are implemented, but for a network of Erlang nodes. A registered name is an alias for a process identity `Pid`. The system monitors globally registered Pids. If a process terminates, the name will also be globally unregistered.

The registered names are stored in replica global name tables on every node. There is no central storage point. Thus, the translation of a name to a `Pid` is extremely quick because it is never a network operation. When any action is taken which results in a change to the global name table all tables on other nodes are automatically updated.

Global locks have lock identities and are set on a specific resource. For instance, the specified resource could be a `Pid` of a process. When a global lock is set access to the locked resource is denied for all other resources other than the lock requester.

Both the registration and lock functionalities are atomic. All nodes involved in these actions will have the same view of the information.

The server also performs the critical task of continuously monitoring changes in node configuration, if a node which runs a globally registered process goes down, the name will be globally unregistered. The server will also maintain a fully connected network. For example, if node `N1` connects to node `N2` (which is already connected to `N3`), the `global` server on `N1` then `N3` will make sure that also `N1` and `N3` are connected. If this is not desired, the command line flag `-connect_all false` must be passed to `init` at boot time. In this case, the name registration facility cannot be used (but the lock mechanism will still work.)

## Exports

```
del_lock(Id)
del_lock(Id, Nodes) -> void()
```

Types:

- `Id = {ResourceId, LockRequesterId}`

- ResourceId = term()
- LockRequesterId = term()
- Nodes = [node()]

Deletes the lock Id synchronously.

`notify_all_name(Name, Pid1, Pid2) -> none`

This function can be used as a name resolving function for `register_name/3` and `re_register_name/3`. It unregisters both Pids, and sends the message `{global_name_conflict, Name, OtherPid}` to both processes.

`random_exit_name(Name, Pid1, Pid2) -> Pid1 | Pid2`

This function can be used as a name resolving function for `register_name/3` and `re_register_name/3`. It randomly chooses one of the Pids for registration and kills the other one.

`random_notify_name(Name, Pid1, Pid2) -> Pid1 | Pid2`

This function can be used as a name resolving function for `register_name/3` and `re_register_name/3`. It randomly chooses one of the Pids for registration, and sends the message `{global_name_conflict, Name}` to the other Pid.

`register_name(Name, Pid)`

`register_name(Name, Pid, Resolve) -> yes | no`

Types:

- Name = term()
- Pid = Pid()
- Resolve = {M, F} where `M:F(Name, Pid, Pid2) -> Pid | Pid2 | none`

Globally notifies all nodes of a new global name in a network of Erlang nodes.

When new nodes are added to the network, they are informed of the globally registered names that already exist. The network is also informed of any global names in newly connected nodes. If any name clashes are discovered, the `Resolve` function is called. Its purpose is to decide which Pid is correct. This function blocks the global name server during its execution. If the function crashes, or returns anything other than one of the Pids, the name is unregistered. This function is called once for each name clash.

There are three pre-defined resolve functions, `random_exit_name`, `random_notify_name` and `notify_all_name`. If no `Resolve` function is defined, `random_exit_name` is used. This means that one of the two registered processes will be selected as correct while the other is killed.

This function is completely synchronous. This means that when this function returns, the name is either registered on all nodes or none.

The function returns `yes` if successful, `no` if it fails. For example, `no` is returned if an attempt is made to register a process with a name that is already in use.

If a process with a registered name dies, or the node goes down, the name is unregistered on all nodes.

`registered_names() -> [Name]`

Types:

- Name = term()

Returns a lists of all globally registered names.

```
re_register_name(Name, Pid)
```

```
re_register_name(Name, Pid, Resolve) -> void()
```

Types:

- Name = term()
- Pid = Pid()
- Resolve = {M, F} where M:F(Name, Pid, Pid2) -> Pid | Pid2 | none

Atomically changes the registered name Name on all nodes to refer to Pid.

The Resolve function has the same behavior as in register\_name.

```
send(Name, Msg) -> Pid
```

Types:

- Name = term()
- Msg = term()
- Pid = Pid()

Sends the message Msg to the globally registered process Name. If Name is not a globally registered name, the calling function will exit with reason {badarg, {Name, Msg}}.

```
set_lock(Id)
```

```
set_lock(Id, Nodes)
```

```
set_lock(Id, Nodes, Retries) -> boolean()
```

Types:

- Id = {ResourceId, LockRequesterId}
- ResourceId = term()
- LockRequesterId = term()
- Nodes = [node()]
- Retries = int() > 0 | infinity

Sets a lock on the specified nodes (or on all nodes if none are specified) on ResourceId for LockRequesterId. If a lock already exists on ResourceId for another requester than LockRequesterId, and Retries is not equal to 0, the process sleeps for a while and will try to execute the action later. When Retries attempts have been made, false is returned, otherwise true. If Retries is infinity, true is eventually returned (unless the lock is never released).

If no value for Retries is given, infinity is used.

This function is completely synchronous.

If a process which holds a lock dies, or the node goes down, the locks held by the process are deleted.

global keeps track of all processes sharing the same lock, i.e. if two processes set the same lock both processes must delete the lock.

This function does not address the problem of a deadlock. A deadlock can never occur as long as processes only lock one resource at a time. But if some processes try to lock two or more resources, a deadlock may occur. It is up to the application to detect and rectify a deadlock.

`start()`

`start_link() -> {ok, Pid} | {error, Reason}`

This function starts the global name server. Normally, the server is started automatically.

`stop() -> void()`

Stops the global name server.

`sync() -> void()`

Synchronizes the global name server with all nodes known to this node. These are the nodes which are returned from `erlang:nodes()`. When this function returns, the global name server will receive global information from all nodes. This function can be called when new nodes are added to the network.

`trans(Id, Fun)`

`trans(Id, Fun, Nodes)`

`trans(Id, Fun, Nodes, Retries) -> Res | aborted`

Types:

- `Id = {ResourceId, LockRequesterId}`
- `ResourceId = term()`
- `LockRequesterId = term()`
- `Fun = fun() | {M, F}`
- `Nodes = [node()]`
- `Retries = int() > 0 | infinity`
- `Res = term()`

Sets a lock on `Id` (using `set_lock/3`). Evaluates `Res = Fun()` if successfully locked and returns `Res`. Returns `aborted` if the lock attempt failed. If `Retries` is set to `infinity`, the transaction will not abort.

`infinity` is the default setting and will be used if no value is given for `Retries`.

`unregister_name(Name) -> void()`

Types:

- `Name = term()`

Globally removes `Name` from the network of Erlang nodes.

`whereis_name(Name) -> Pid() | undefined`

Types:

- `Name = term()`

Returns either an atom `undefined`, or a `Pid` which is globally associated with `Name`.

# global\_group (Module)

The global group function makes it possible to group the nodes in a system into partitions, each partition having its own global name space, refer to `global(3)`. These partitions are called global groups.

The main advantage of dividing systems to global groups is that the background load decreases while the number of nodes to be updated is reduced when manipulating globally registered names.

The `global_groups`-key in the `.config` file defines the global groups:

```
{global_groups, [{GroupName, [Node]}] }
```

`GroupName` is an `atom()` naming a global group.

`Node` is an `atom()` naming a node.

The command `erl -config File` starts a node with a configuration file named `File.config`. If the `global_groups`-key is not defined the system will start as a whole, without any partitions. When the key is not defined, the services of this function will not give any extra value to `global(3)`.

For the processes and nodes to run smoothly using this function the following criteria must be met:

- The global group function must have a server process, `global_group`, running on each node.  
*NOTE:* The processes are automatically started and synchronized when a node is started.
- All processes must agree with the group definition in the immediate global group. If two nodes do not agree, these nodes will not synchronize their name space and an error message will be logged in the error logger.  
Example: If one node has an illegal global group definition, such a node will run isolated from the other nodes regarding the global name space; but not regarding other system functions, e.g distribution of applications, refer to chapter *NOTE* below.
- Nodes can only belong to one global group.

When the global group definitions are to be changed in a system upgrade, the upgrade must fulfill the following steps:

1. First, all nodes which are to be removed from a global group must be taken down.
2. Nodes which are not affected by the redefinition of the global groups are to be upgraded to be aware of the new global group definitions.  
*NOTE:* All nodes in the system, also nodes in unchanged global groups, must be upgraded. This because e.g `send` must have an accurate view of the total system.
3. Finally, all nodes which are new to a global group can be started.

When a non partitioned system is to be upgraded to become a partitioned system, all nodes belonging to a global group will be disconnected from all nodes not belonging to its immediate global group.

## Exports

`global_groups()` -> {OwnGroupName, [OtherGroupName]} | undefined

Types:

- OwnGroupName = atom()
- OtherGroupName = atom()
- ErrorMsg = term()

Returns the names of all the global groups known to the immediate global group.

`info()` -> [{state, State}, {own\_group\_name, atom()}, {own\_group\_nodes, [Node]}, {synced\_nodes, [Node]}, {sync\_error, [Node]}, {no\_contact, [Node]}, {other\_groups, Other\_grps}, {monitoring, [pid()]}]

Types:

- State = no\_conf | synced
- Other\_grps = [{OtherGrpName, [Node]}]
- OtherGrpName = atom()
- Node = atom()

Returns the state of the global group process. In the following 'nodes' refers to nodes in the immediate global group. `synced_nodes` lists the nodes this node is synchronized with at this moment. `own_group_nodes` lists the nodes defining the own global group. `sync_error` lists the nodes with this node could not be synchronize. `no_contact` lists nodes with this node do not yet have established contact. `other_groups` shows the definition of the other global groups in the system. `monitoring` lists the processes which have subscribed on `nodeup` and `nodedown` messages.

`monitor_nodes(Flag)` -> ok

Types:

- Flag = bool()

The requesting process receives {`nodeup`, Node} and {`nodedown`, Node} messages about the nodes from the immediate global group. If the flag `Flag` is set to `true` the service is turned on; `false` turns it off.

`own_nodes()` -> [Node] | {error, ErrorMsg}

Types:

- Node = atom()
- ErrorMsg = term()

Returns the names of all nodes from the immediate global group, despite of the status of the nodes. Use `info/0` to get the information of the current status of the nodes.

`registered_names({node, Node})` -> [Name] | {error, ErrorMsg}

`registered_names({group, GlobalGroupName})` -> [Name]

Types:

- Name = term()

- Node = atom()
- GlobalGroupName = atom()
- ErrorMsg = term()

Returns a lists of all globally registered names on the specified node or from the specified global group.

```
send(Name, Msg) -> Pid | {badarg, Msg} | {error, ErrorMsg}
```

```
send({node, Node}, Name, Msg) -> Pid | {badarg, Msg} | {error, ErrorMsg}
```

```
send({group, GlobalGroupName}, Name, Msg) -> Pid | {badarg, Msg} | {error, ErrorMsg}
```

Types:

- GlobalGroupName = atom()
- Msg = term()
- Name = term()
- Node = atom()
- Pid = pid()
- ErrorMsg = term()

send/2 searches for the registered Name in all global groups defined, in the order of appearance in the .config-file, until the registered name is found or all groups are searched. If Name is found, the message Msg is sent to it. If it is not found, the function exits with reason {badarg, {Name, Msg}}.

send/3 searches for the registered Name in either the specified node or the specified global group. If the registered name is found, the message Msg is sent to that process. If Name is not found, the function exits with reason {badarg, {Name, Msg}}.

```
sync() -> ok
```

sync synchronizes the global name servers on the nodes in the immediate global group. It also unregisters the names registered in the immediate global group on known nodes outside to the immediate global group.

If it the global\_groups definition is invalid, the function exits with reason {error, {'invalid global\_groups definition', NodeGrpDef}}.

```
whereis_name(Name) -> Pid | undefined | {error, ErrorMsg}
```

```
whereis_name({node, Node}, Name) -> Pid | undefined | {error, ErrorMsg}
```

```
whereis_name({group, GlobalGroupName}, Name) -> Pid | undefined | {error, ErrorMsg}
```

Types:

- GlobalGroupName = atom()
- Name = term()
- Node = atom()
- Pid = pid()

whereis\_name/1 searches for the registered Name in all global groups defined, in the order of appearance in the .config-file, until the registered name is found or all groups are searched.

whereis\_name/2 searches for the registered Name in either the specified node or the specified global group.

Returns either the atom undefined, or the Pid which is associated with Name.

`start()`

`start_link() -> {ok, Pid} | {error, Reason}`

This function starts the global group server. Normally, the server is started automatically.

`stop() -> void()`

Stops the global group server.

## NOTE

In the situation where a node has lost its connections to other nodes in its global group but has connections to nodes in other global groups, a request from the other global group may produce an incorrect or misleading result. When this occurs the isolated node may not have accurate information, for example, about registered names in its global group.

Note also that the send function is not secure.

Distribution of applications is highly dependent of the global group definitions. It is not recommended that an application is distributed over several global groups of the obvious reason that the registered names may be moved to another global group at failover/takeover. There is nothing preventing doing this, but the application code must in such case handle the situation.

## SEE ALSO

`erl(1)`, `global(3)`

# heart (Module)

The `heart` module sends periodic heartbeats to an external port program, which is also named `heart`. The purpose of the `heart` port program is to check that the Erlang system it is supervising is still running. If the port program has not received any heartbeats within `HEART_BEAT_TIMEOUT` (default is 60 seconds) from the last one, the system can be rebooted. Also, if the system is equipped with a hardware watchdog timer and is running Solaris, the watchdog can be used to supervise the entire system.

This module is started by the `init` module during system start-up. The `-heart` command line flag determines if the `heart` module should start.

If the system should be rebooted because of missing heart-beats, or a terminated Erlang system, the environment variable `HEART_COMMAND` has to be set before the system is started. If this variable is not set, a warning text will be printed but the system will not reboot. However, if the hardware watchdog is used, it will trigger a reboot `HEART_BEAT_BOOT_DELAY` seconds later nevertheless (default is 60).

To reboot on the `WINDOWS` platform `HEART_COMMAND` can be set to `heart -shutdown` (included in the Erlang delivery) or of course to any other suitable program which can activate a reboot.

The hardware watchdog will not be started under Solaris if the environment variable `HW_WD_DISABLE` is set.

The `HEART_BEAT_TIMEOUT` and `HEART_BEAT_BOOT_DELAY` environment variables can be used to configure the heart timeouts, they can be set in the operating system shell before `erl -heart` is started or can be passed on the command line like this: `erl -heart -env HEART_BEAT_TIMEOUT 30`.

The value (in seconds) must be in the range  $10 < X \leq 65535$ .

It should be noted that if the system clock is adjusted with more than `HEART_BEAT_TIMEOUT` seconds `heart` will timeout and try to reboot the system. This can happen for example if the system clock is adjusted automatically by use of NTP (Network Time Protocol).

## Exports

```
start() -> {ok, Pid} | ignore | {error, What}
```

Types:

- `Pid = pid()`
- `What = void()`

Starts the `heart` program. This function returns `ignore` if the `-heart` command line flag is not supplied.

`set_cmd(Cmd) -> ok | {error, {bad_cmd, Cmd}}`

Types:

- `Cmd = string()`

Sets a temporary reboot command. This command is used if a `HEART_COMMAND` other than the one specified with the environment variable should be used in order to reboot the system. The new Erlang system will (if it misbehaves) use the environment variable `HEART_COMMAND` to reboot.

The length of the `Cmd` command string must be less than 2047 characters.

`clear_cmd() -> ok`

Clears the temporary boot command. If the system terminates, the normal `HEART_COMMAND` is used to reboot.

# inet (Module)

Inet provides access to TCP/IP protocols.

Some functions returns a `hostent` record. Use this line in your module `-include_lib("kernel/include/inet.hrl")`. to include the record definition.

`h_addr_list` List of addresses for this host  
`h_addrtype` Type of address: `inet` or `inet6`  
`h_aliases` List of aliases (additional names for host)  
`h_length` Length of address in bytes  
`h_name` Official name for host

Addresses as inputs to functions can be either a string or a tuple. For instance, the IP address 150.236.20.73 can be passed to `gethostbyaddr/1` either as the string "150.236.20.73" or as the tuple `{150, 236, 20, 73}`. Addresses returned by any function in the `inet` module will be a tuple.

Hostnames may be specified as either atoms or a strings.

Where an address family is required, valid values are `inet` (standard IPv4 addresses) or `inet6` (IPv6).

## Exports

`format_error(Tag)`

Types:

- `Tag = atom()`

Returns a diagnostic error string. See the section below for possible `Tag` values and the corresponding strings.

`gethostbyaddr(Address) -> {ok, Hostent} | {error, Reason}`

Types:

- `Address = address()`
- `Hostent = hostent()`

Returns a `hostent` record given an address.

`gethostbyname(Name) -> {ok, Hostent} | {error, Reason}`

Types:

- Hostname = hostname()
- Hostent = hostent()

Returns a `hostent` record given a hostname.

`gethostbyname(Name, Family) -> {ok, Hostent} | {error, Reason}`

Types:

- Hostname = hostname()
- Family = family()
- Hostent = hostent()

Returns a `hostent` record given a hostname, restricted to the given address family.

`gethostname() -> {ok, Name} | {error, Reason}`

Types:

- Hostname = hostname()

Returns the local hostname. Will never fail.

`sockname(Socket) -> {ok, {IP, Port}} | {error, Reason}`

Types:

- Socket = socket()
- Address = address()
- Port = integer()

Returns the local address and port number for a socket.

`peername(Socket) -> {ok, {Address, Port}} | {error, Reason}`

Types:

- Socket = socket()
- Address = address()
- Port = integer()

Returns the address and port for the other end of a connection.

`port(Socket) -> {ok, Number}`

Types:

- Socket = socket()
- Number = integer()

Returns the local port number for a socket.

`close(Socket) -> ok`

Types:

- Socket = socket()

Closes a socket of any type.

`getaddr(IP, inet) -> {ok, {A1, A2, A3, A4}} | {error, Reason}`

Types:

- IP = {A1, A2, A3, A4} | string() | atom()
- A1 = A2 = A3 = A4 = integer()
- Reason = term()

Returns the IP-address as a tuple with integers for IP which can be an IP-address a single hostname or a fully qualified hostname. At present only IPv4 addresses (the `inet` argument) is supported, but the function is prepared to support IPv6 (`inet6`) in a near future.

`setopts(Socket, Options) -> ok | {error, Reason}`

Types:

- Socket = term()
- Options = list()

Sets one or more options for a socket. The following options are available:

`{active, Boolean}` If the `active` option is `true`, which is the default, everything received from the socket will be sent as messages to the receiving process. If the `active` option is set to `false` (passive mode), the process must explicitly receive incoming data by calling `gen_tcp:recv/N` or `gen_udp:recv/N` (depending on the type of socket). Note: Passive mode provides flow control; the other side will not be able send faster than the receiver can read. Active mode provides no flow control; a fast sender could easily overflow the receiver with incoming messages. Use active mode only if your high-level protocol provides its own flow control (for instance, acknowledging received messages) or the amount of data exchanged is small.

`{broadcast, Boolean}` Enable/disable permission to send broadcasts. (UDP)

`{dontroute, true|false}` Use `{dontroute, true}` to enable/disable routing bypass for outgoing messages.

`{header, Size}` This option is only meaningful if the `binary` option was specified when the socket was created. If the `header` option is specified, the first `Size` number bytes of data received from the socket will be elements of a list, and the rest of the data will be a binary given as the tail of the same list. If for example `Size=2` the data received will match `[Byte1, Byte2|Binary]`.

`{keepalive, Boolean}` (TCP/IP sockets) Enables periodic transmission on a connected socket, when no other data is being exchanged. If the other end does not respond, the connection is considered broken and an error message will be sent to the controlling process. Default disabled.

`{nodelay, Boolean}` If `Boolean` is `true`, the `TCP_NODELAY` option is turned on for the socket, which means that even small amounts of data will be sent immediately. (TCP/IP)

`{packet, PacketType}` Defines the type of packets to use for a socket. The following values are valid:

`raw` | `0` No packaging is done.

1 | 2 | 4 Packets consist of a header specifying the number of bytes in the packet, followed by that number of bytes. The length of header can be one, two, or four bytes; the order of the bytes is big-endian. Each send operation will generate the header, and the header will be stripped off on each receive operation.

asn1 | cdr | sunrm | fcgi These packet types only have effect on receiving. When sending a packet, it is the responsibility of the application to supply a correct header. On receiving, however, there will be one message sent to the controlling process for each complete packet received, and, similarly, each call to `gen_tcp:recv/N` returns one complete packet. The header is *not* stripped off. The meaning of the packet types is as follows:

asn1 - ASN.1 BER,  
 sunrm - Sun's RPC encoding,  
 cdr - Corba,  
 fcgi - Fast CGI.

{`recbuf`, Integer} Gives the size of the receive buffer to use for the socket.

{`reuseaddr`, Boolean} Allows or disallows local reuse of port numbers. By default, reuse is disallowed.

{`sndbuf`, Integer} Gives the size of the send buffer to use for the socket.

## ERRORS

The possible error reasons and the corresponding diagnostic strings returned by `format_error/1` are as follows:

e2big argument list too long  
 eaccess permission denied  
 eaddrinuse address already in use  
 eaddrnotavail can't assign requested address  
 eadv advertise error  
 eafnosupport address family not supported by protocol family  
 eagain resource temporarily unavailable  
 ealign EALIGN  
 ealready operation already in progress  
 ebade bad exchange descriptor  
 ebadf bad file number  
 ebadfd file descriptor in bad state  
 ebadmsg not a data message  
 ebadr bad request descriptor  
 ebadrpc RPC structure is bad  
 ebadrqc bad request code  
 ebads1t invalid slot

ebfont bad font file format  
ebusy file busy  
echild no children  
echnrg channel number out of range  
ecomm communication error on send  
econnaborted software caused connection abort  
econnrefused connection refused  
econnreset connection reset by peer  
edeadlk resource deadlock avoided  
edeadlock resource deadlock avoided  
edestaddrreq destination address required  
edirty mounting a dirty fs w/o force  
edom math argument out of range  
edotdot cross mount point  
edquot disk quota exceeded  
eduppkg duplicate package name  
eexist file already exists  
efault bad address in system call argument  
efbig file too large  
ehostdown host is down  
ehostunreach host is unreachable  
eidrm identifier removed  
einit initialization error  
einprogress operation now in progress  
eintr interrupted system call  
einval invalid argument  
eio I/O error  
eisconn socket is already connected  
eisdir illegal operation on a directory  
eisnam is a named file  
el2hlt level 2 halted  
el2nsync level 2 not synchronized  
el3hlt level 3 halted  
el3rst level 3 reset  
elbin ELBIN  
elibacc can not access a needed shared library  
elibbad accessing a corrupted shared library  
elibexec can not exec a shared library directly  
elibmax attempting to link in more shared libraries than system limit  
elibscn .lib section in a.out corrupted  
elnrng link number out of range

eLOOP too many levels of symbolic links  
emfile too many open files  
mlink too many links  
emsgsize message too long  
emultihop multihop attempted  
enametoolong file name too long  
enavail not available  
enet ENET  
enetdown network is down  
enetreset network dropped connection on reset  
enetunreach network is unreachable  
enfile file table overflow  
enoano anode table overflow  
enobufs no buffer space available  
enocsi no CSI structure available  
enodata no data available  
enodev no such device  
enoent no such file or directory  
enoexec exec format error  
enolck no locks available  
enolink link has be severed  
enomem not enough memory  
enomsg no message of desired type  
enonet machine is not on the network  
enopkg package not installed  
enoprotoopt bad proocol option  
enospc no space left on device  
enosr out of stream resources or not a stream device  
enosym unresolved symbol name  
enosys function not implemented  
enotblk block device required  
enotconn socket is not connected  
enotdir not a directory  
enotempty directory not empty  
enotnam not a named file  
enotsock socket operation on non-socket  
enotsup operation not supported  
enotty inappropriate device for ioctl  
enotuniq name not unique on network  
enxio no such device or address  
eopnotsupp operation not supported on socket

---

eperm not owner  
epfnosupport protocol family not supported  
epipe broken pipe  
eproclim too many processes  
eprocunavail bad procedure for program  
eprogmismatch program version wrong  
eprogunavail RPC program not available  
eproto protocol error  
eprotonosupport protocol not supported  
eprototype protocol wrong type for socket  
erange math result unrepresentable  
erefused EREFUSED  
eremchg remote address changed  
eremdev remote device  
eremote pathname hit remote file system  
eremoteio remote i/o error  
eremoterelease EREMOTERELEASE  
erofs read-only file system  
erpcmismatch RPC version is wrong  
erremote object is remote  
eshutdown can't send after socket shutdown  
esocktnosupport socket type not supported  
espipe invalid seek  
esrch no such process  
esrmnt srmount error  
estale stale remote file handle  
esuccess Error 0  
etime timer expired  
etimedout connection timed out  
etoomanyrefs too many references  
etxtbsy text file or pseudo-device busy  
euclean structure needs cleaning  
eunatch protocol driver not attached  
eusers too many users  
eversion version mismatch  
ewouldblock operation would block  
exdev cross-domain link  
exfull message tables full  
nxdomain the hostname or domain name could not be found

# init (Module)

`init` is pre-loaded into the system before the system starts and it coordinates the start-up of the system. The first function evaluated at start-up is `boot(Bootargs)`, where `Bootargs` is a list of the arguments supplied to the Erlang system from the local operating system. The Erlang code for the module `init` is always pre-loaded.

`init` reads a boot script which contains instructions on how to initiate the system. The default boot script (`start.boot`) is in the directory `<ERL_INSTALL_DIR>/bin`.

`init` contains functions to fetch command line flags, or arguments, supplied to the Erlang system.

`init` also contains functions to restart, reboot, and stop the system.

## Exports

`boot(BootArgs) -> void()`

Types:

- `BootArgs = [term()]`

Erlang is started with the command `erl <script-flags> <user-flags>`.

`erl` is the name of the Erlang start-up script. `<script-flags>`, described in `erl(1)`, are read by the script. `<user-flags>` are put into a list and passed as `Args` to `boot/1`.

The `boot/1` function interprets the `boot`, `mode`, and `s` flags. These are described in `COMMAND LINE FLAGS`.

If the `boot` function finds other arguments starting with the character `-`, that argument is interpreted as a flag with zero or more values. It ends the previous argument. For example:

```
erl -s foo bar -charles peterson
```

This starts the Erlang system, evaluates `foo:bar()`, and sets the flag `-charles`, which has the associated value `peterson`.

Other arguments which are passed to the `boot` function, and do not fit into the above description, are passed to the `init` loop as plain arguments.

The special flag `--` can be used to separate plain arguments to `boot`.

`get_arguments() -> Flags`

Types:

- `Flags = [{Flag,[Value]}]`
- `Flag = atom()`

- Value = string()

Returns all flags given to the system.

`get_argument(Flag) -> {ok, Values} | error`

Types:

- Flag = atom()
- Values = [FValue]
- FValue = [Value]
- Value = string()

Returns all values associated with `Flag`. If `Flag` is provided several times, each `FValue` is returned in preserved order.

`get_args() -> [Arg]`

Types:

- Arg = atom()

Returns the additional plain arguments as a list (possibly empty).

`restart() -> void()`

The system is restarted inside the running Erlang node, which means that the emulator is not restarted. All applications are taken down smoothly, all code is unloaded, and all ports are closed before the system is booted again in the same way as initially started. The same `BootArgs` are used again.

To limit the shutdown time, the time `init` is allowed to spend taking down applications, the `-shutdown_time` command line flag should be used.

`reboot() -> void()`

All applications are taken down smoothly, all code is unloaded, and all ports are closed before the Erlang node terminates. If the `-heart` system flag was given, the `heart` program will try to reboot the system. Refer to the `heart` module for more information.

In order to limit the shutdown time, the time `init` is allowed to spend taking down applications, the `-shutdown_time` command line flag should be used.

`stop() -> void()`

All applications are taken down smoothly, all code is unloaded, and all ports are closed before the system terminates. If the `-heart` system flag was given, the `heart` program is terminated before the Erlang node terminates. Refer to the `heart` module for more information.

In order to limit the shutdown time, the time `init` is allowed to spend taking down applications, the `-shutdown_time` command line flag should be used.

`get_status() -> {InternalStatus, ProvidedStatus}`

Types:

- InternalStatus = starting | started | stopping
- ProvidedStatus = term()

The current status of the `init` process can be inspected. During system start (initialization), `InternalStatus` is `starting`, and `ProvidedStatus` indicates how long the boot script has been interpreted. Each `{progress, Info}` term interpreted in the boot script affects the `ProvidedStatus` status, i.e., `ProvidedStatus` gets the value of `Info`.

`script_id()` -> `Id`

Types:

- `Id = term()`

Get the identity of the boot script used to boot the system. `Id` can be any Erlang term. In the delivered boot scripts, `Id` is `{Name, Vsn}`. `Name` and `Vsn` are strings.

## Command Line Flags

The `init` module interprets the following flags:

- boot File** Specifies the name of the boot script, `File.boot`, used to start the system. Unless `File` contains an absolute path, the system searches for `File.boot` in the current and `<ERL_INSTALL_DIR>/bin` directories. If this flag is omitted, the `<ERL_INSTALL_DIR>/bin/start.boot` boot script is used.
- boot\_var Var Directory [Var Directory]** If the boot script used contains another path variable than `$ROOT`, that variable must have a value assigned in order to start the system. A boot variable is used if user applications are installed in a different location than underneath the `<ERL_INSTALL_DIR>/lib` directory. `$Var` is expanded to `Directory` in the boot script.
- mode Mode** The mode flag indicates if the system will load code automatically at runtime, or if all code should be loaded during system initialization. `Mode` can be either `interactive` (allow automatic code loading) or `embedded` (load all code during start-up).
- shutdown\_time Time** Specifies how long time (in ms) the `init` process is allowed to spend shutting down the system. If `Time` milliseconds has elapsed, all processes still existing are *killed*. If `-shutdown_time` is not specified, the default time is `infinity`.
- s Module [Function [Args]]** Evaluate the function during system initialization. `Function` defaults to `start` and `Args` to `[]`. If the function call ends abnormally, the Erlang system stops with an error message. The arguments after `-s` are used as arguments to Erlang functions. All arguments are passed as atoms. For example:  

```
erl -s foo -s foo bar -s foo bar baz 1 2
```

This starts the Erlang system and then evaluates the following Erlang functions:

```
foo:start()
foo:bar()
foo:bar([baz, '1', '2']).
```

The functions are executed sequentially in the initialization process, which then terminates normally and passes control to the user. This means that a `-s` call which doesn't terminate will block further processing; to avoid this, use some variant of `spawn` in such cases.

**-init\_debug** The `init` process writes some debug information while interpreting the boot script.

## Example

```
erl -- a b -children thomas claire -ages 7 3 -- x y
1> init:get_args().
[a, b, x, y]
2> init:get_argument(children).
{ok, [{"thomas", "claire"]}]
3> init:get_argument(ages).
{ok, [{"7", "3"}]}
4> init:get_argument(silly).
error
```

## See also

`erl_prim_loader(3)`, `heart(3)`

# net\_adm (Module)

This module contains various network utility functions.

## Exports

`host_file()`

This function reads the `.hosts.erlang` file. It returns the hosts in this file as a list, or it returns `{error, Reason}` if the file cannot be found.

`dns_hostname(Host)`

This function calls `epmd` for the fully qualified name (DNS) of `Host`. It returns `{ok, Longhostname}` if the call is successful, or `{error, Host}` if `Host` cannot be located by DNS.

`localhost()`

This function returns the fully qualified name of the local host, if it can be found by DNS.

`names()`, `names(Host)`

This function returns `{ok, List}` or `{error, Reason}`. `List` is a list of tuples on the form `{Name, Port}`. For example: `net_adm:names(elrond) -> {ok, [{"foo", 61178}, {"ts", 61160}]}`.

`ping(Node)`

This function tries to set up a connection to `Node`. It returns `pang` if it fails, and `pong` if it is successful.

`world ()`, `world (verbose)`

This function runs `epmd - names` on all hosts which are specified in the Erlang host file `.hosts.erlang`, collects the replies and then evaluates `ping` on all those nodes. Accordingly, connections are created to all nodes which are running on the hosts specified in the `.hosts.erlang` file. An error message is printed if another user node is found when this is done.

This function can be useful when a node is started, but the names of the other nodes in the network are not initially known.

`world_list (Hostlist)`, `world_list (Hostlist, verbose)`

These functions are the same as `world/0` and `world/1`, but instead of reading the hostfile from `.hosts.erlang` the hosts are specified in `Hostlist`.

## Files

The `.hosts.erlang` file consists of a number of host names written as Erlang terms. It can be located in the current work directory, `$HOME/.hosts.erlang`, or `code:root_dir()/.hosts.erlang`. The format of the `.hosts.erlang` file must be one host name per line. The host names must be within quotes as shown in the following examples:

```
'super.eua.ericsson.se'.  
'renat.eua.ericsson.se'.  
'grouse.eua.ericsson.se'.  
'gauffin1.eua.ericsson.se'.  
^ (new line)
```

# net\_kernel (Module)

The net kernel is a system process which must be running for distributed Erlang to work. The purpose of this process is to implement parts of the BIFs `spawn/4` and `spawn_link/4`, and to provide authentication and monitoring of the network.

An Erlang system can be started from the UNIX command line as follows:

```
% erl -name foobar
```

With this command line, the `net_kernel` is started as `net_kernel:start([foobar])`.

This is done by the system itself, but the `start([Name])` function can also be called directly from the normal Erlang shell prompt, and a normal Erlang system is then converted to a node. The kernel can be shut down with the function `stop()`, but only if the kernel was not started by the system itself. The node is then converted into a normal Erlang system. All other nodes on the network will regard this as a total node crash.

If the system is started as `% erl -sname foobar`, the node name of the node will be `foobar@Host`, where `Host` is the short name of the host (not the fully qualified domain name). The `-name` flag gives a node with the fully qualified domain name.

## Exports

`kernel_apply(M, F, A)`

As the net kernel runs in the 'user space', it is easy to provide another net kernel which is tailor made for a specific application. For example, the user supplied kernel can limit the set of registered processes which can be accessed from remote nodes, or it can spawn a new process for each `{nodeup, Node}` message and perform some application specific user authentication, a log-in procedure for example. The `kernel_apply(M, F, A)` function is supplied for this purpose.

`monitor_nodes(Flag)`

A process which evaluates this function receives copies of the `{nodeup, Node}` and `{nodedown, Node}` messages that the net kernel receives from the runtime system. The flag `Flag` is set to `true` to turn the service on, and `false` to turn it off.

`allow(NodeList)`

In a simple way, this function limits access to a node from a specific number of named nodes. A node which evaluates this function can only be accessed from nodes listed in the `NodeList` variable. Any access attempts made from nodes not listed in `NodeList` are rejected.

# os (Module)

The functions in this module are operating system specific. Careless use of these functions will result in programs that will only run on a specific platform. On the other hand, with careful use these functions can be of help in enabling a program to run on most platforms.

## Exports

`cmd(Command) -> string()`

Types:

- `Command = string() | atom()`

Executes `Command` in a command shell of the target OS and returns the result as a string. This function is a replacement of the previous `unix:cmd/1`; on a Unix platform they are equivalent.

Examples:

```
LsOut = os:cmd("ls"), % on unix platform
```

```
DirOut = os:cmd("dir"), % on Win32 platform
```

`find_executable(Name) -> Filename | false`

`find_executable(Name, Path) -> Filename | false`

Types:

- `Name = string()`
- `Path = string()`
- `Filename = string()`

These two functions look up an executable program given its name and a search path, in the same way as the underlying operating system. `find_executable/1` uses the current execution path (i.e., the environment variable `PATH` on Unix and Windows). `Path`, if given, should conform to the syntax of execution paths on the operating system. The absolute filename of the executable program `Name` is returned, or `false` if the program was not found.

`getenv(VarName) -> Value | false`

Types:

- `Varname = string()`
- `Value = string()`

Returns the `Value` of the environment variable `VarName`, or `false` if the environment variable is undefined.

`getpid() -> Value`

Types:

- `Value = string()`

Returns the process identifier of the current Erlang emulator in the format most commonly used by the operating system environment. `Value` is returned as a string containing the (usually) numerical identifier for a process. On Unix, this is typically the return value of the `getpid()` system call. On VxWorks, `Value` contains the task id (decimal notation) of the erlang task. On Windows, the process id as returned by the `GetCurrentProcessId()` system call is used.

`type() -> {Osfamily, Osname} | Osfamily`

Types:

- `Osfamily = atom() = win32 | unix | vxworks`
- `Osname = atom()`

Returns the `Osfamily` and, in some cases, `Osname` of the current operating system.

On Unix, `Osname` will be same string that `uname -s` returns, but in lower case. For instance, on Solaris 1 and 2, the atom `sunos` will be returned.

In Windows, `Osname` will be either `nt` (on Windows NT), or `windows` (on Windows 95). On VxWorks `Osfamily` alone is returned, i.e. the atom `vxworks`.

**Note:**

Think twice before using this function. Use the `filename` module if you want to inspect or build file names in a portable way. Avoid matching on the `Osname` atom.

`version() -> {Major, Minor, Release} | VersionString`

Types:

- `Major = Minor = Release = integer()`
- `VersionString = string()`

Returns the operating system version. On most systems, this function returns a tuple, but a string will be returned instead if the system has versions which cannot be expressed as three numbers.

**Note:**

Think twice before using this function. If you still need to use it, always call `os:type()` first.

# pg2 (Module)

This module implements process groups. The groups in this module differ from the groups in the module `pg` in several ways. In `pg`, each message is sent to all members in the group. In this module, each message may be sent to one, some, or all members.

A group of processes can be accessed by a common name. For example, if there is a group named `foobar`, there can be a set of processes (which can be located on different nodes) which are all members of the group `foobar`. There is no special functions for sending a message to the group. Instead, client functions should be written with the functions `get_members/1` and `get_local_members/1` to find out which process are members of the group. Then the message can be sent to one or more members of the group.

If a member terminates, it is automatically removed from the group.

## Exports

`create(Name) -> void()`

Types:

- `Name = term()`

Creates a new, empty process group. The group is globally visible on all nodes. If the group exists, nothing happens.

`delete(Name) -> void()`

Types:

- `Name = term()`

Deletes a process group.

`get_closest_pid(Name) -> Pid | {error, Reason}`

Types:

- `Name = term()`

This is a useful dispatch function which can be used from client functions. It returns a process on the local node, if such a process exist. Otherwise, it chooses one randomly.

`get_members(Name) -> [Pid] | {error, Reason}`

Types:

- `Name = term()`

Returns all processes in the group `Name`. This function should be used from within a client function that accesses the group. It is then optimized for speed.

```
get_local_members(Name) -> [Pid] | {error, Reason}
```

Types:

- `Name = term()`

Returns all processes running on the local node in the group `Name`. This function should be used from within a client function that accesses the group. It is then optimized for speed.

```
join(Name, Pid) -> ok | {error, Reason}
```

Types:

- `Name = term()`

Joins the process `Pid` to the group `Name`.

```
leave(Name, Pid) -> ok | {error, Reason}
```

Types:

- `Name = term()`

Makes the process `Pid` leave the group `Name`.

```
which_groups() -> [Name]
```

Types:

- `Name = term()`

Returns a list of all known groups.

```
start()
```

```
start_link() -> {ok, Pid} | {error, Reason}
```

Starts the `pg2` server. Normally, the server does not need to be started explicitly, as it is started dynamically if it is needed. This is useful during development, but in a target system the server should be started explicitly. Use configuration parameters for `kernel` for this.

## See Also

`kernel(3)`, `pg(3)`

# rpc (Module)

This module contains services which are similar to remote procedure calls. It also contains broadcast facilities and parallel evaluators. A remote procedure call is a method to call a function on a remote node and collect the answer. It is used for collecting information on a remote node, or for running a function with some specific side effects on the remote node.

## Exports

`start()`

Starts the rpc server. Normally, this is not necessary because the rpc server is started automatically.

`stop()`

Stops the rpc server.

`call(Node, Module, Function, Args)`

Evaluates `apply(Mod, Fun, Args)` on the node `Node` and returns a value, or `{badrpc, Reason}` if the call fails.

`cast(Node, Module, Function, Args)`

Causes the expression `apply(Mod, Fun, Args)` to be evaluated on `Node`. No response is delivered and the process which makes the call is not suspended until the evaluation is complete, as is the case with `call/4`. The function immediately returns `true`. Example:

```
> rpc:cast(Node, erlang, halt, [])
```

This function shuts down the node `Node`.

The following function also shuts down the node, but the call returns the tuple `{badrpc, noconnection}`

```
> rpc:call(Node, erlang, halt, [])
```

`block_call(Node, Mod, Fun, Args)`

The `call/4` function causes the server at `Node` to create a new process for each request. This means that several RPCs can be active concurrently. The `rpc` server is not affected if a request does not return a value. This function can be used if the intention of the call is to block the `rpc` server from any other incoming requests until the request has been handled. The function can also be used for efficiency reasons when very small fast functions are evaluated, for example BIFs that are guaranteed not to suspend.

```
> rpc:block_call(Node, erlang, whereis, [file_server]),
```

Returns the Pid of the file server at `Node`.

```
server_call(Node, Name, ReplyWrapper, Msg)
```

This function is used when interacting with a server called `Name` at node `Node`. It is assumed that the server receives messages in the format `{From, Request}` and replies in the format `From ! {ReplyWrapper, node(), Reply}`. This function makes such a server call and ensures that the entire call is packed into an atomic transaction which either succeeds or fails. It never hangs, unless the server itself hangs.

The function, for example, is used by the `call/4` function to interact with the `rpc` server at a remote node. The function returns `{error, Reason}`, or the answer as produced by the server `Name`.

```
abcast(Name, Mess)
```

Broadcasts the message `Mess` asynchronously to the registered process `Name` on all nodes, including the current node.

```
abcast(Nodes, Name, Mess)
```

The same as `abcast/2`, but only to the nodes `Nodes`.

```
sbcast(Name, Msg)
```

Broadcasts to all nodes synchronously and returns a list of the nodes which have `Name` as a registered server. Returns `{Goodnodes, Badnodes}`.

It is synchronous in the sense that it is known that all servers have received the message when we return from the call. It is not possible to know that the servers have actually processed the message.

Any further messages sent to the servers, after this function has returned, will be received by all servers after this message .

```
sbcast(Nodes, Name, Msg)
```

As `sbcast/2` but only to the nodes in `Nodes`.

```
eval_everywhere(Mod, Fun, Args)
```

Evaluates the expression `apply(Mod, Fun, Args)` on all nodes. No answers are collected.

```
eval_everywhere(Nodes, Mod, Fun, Args)
```

Evaluates the expression `apply(Mod, Fun, Args)` on the nodes `Nodes`.

```
multicall(M, F, A)
```

In contrast to an RPC, a multicall is an RPC which is sent concurrently from one client to multiple servers. This is useful for collecting some information from a set of nodes, or for calling a function on a set of nodes to achieve some side effects. It is semantically the same as iteratively making a series of RPCs on all the nodes, but the multicall is faster as all the requests are sent at the same time and are collected one by one as they come back.

The function `multicall/3` evaluates the expression `apply(M, F, A)` on all nodes and collects the answers. It returns `{Replies, Badnodes}`, where `Badnodes` is a list of the nodes that terminated during computation and `Replies` is a list of the return values. This is useful when new object code is to be loaded on all nodes in the network.

```
%% Find object code for module Mod
{Mod, File, Bin} = code:get_object_code(Mod),

%% and load it on all nodes including this one
{Replies, _} = rpc:multicall(code, load_binary, [Mod, File, Bin]),

%% and then maybe check the Replies list.
```

This is an example of the side effects the RPCs may produce.

```
multicall(Nodes, M, F, A)
```

Executes the multicall only on the nodes `Nodes`.

```
multi_server_call(Name, Msg)
```

This function is used by the `multicall/3` function when it interacts with the `rpc` server. The function can, of course, also be used for interacting with other servers.

The function sends `Msg` to `Name` on all nodes, and collects the answers. It returns `.br {Replies, Badnodes}`, where `Badnodes` is a list of the nodes which failed during the call. This function assumes that if a request sent to a server called `Name`, the server replies in the form `{Name, node(), Reply}`. Otherwise, the function will hang. It also assumes that the server receives messages in the form `{From, Msg}`, and then replies as `.br From ! {Name, node(), Reply}`.

The function also hangs if any one of the servers involved crash, or do not exist, during the call. This function, then, assumes that the servers are stable. If any of the nodes involved crash, the function still works and the crashed node appears in the `Badnodes` list of the return value.

The replies are not ordered in any particular way.

```
multi_server_call(Nodes, Name, Msg)
```

The same as above, but `Msg` is only sent to `Nodes`.

```
safe_multi_server_call(Name, Msg)
```

The same as the `multi_server_call/2`, except that this function handles the case where the remote node exists, but no server called `Name` exists there. This call is also slightly slower than `multi_server_call/2` since all request go via the `rpc` server at the remote sites.

```
safe_multi_server_call(Nodes, Name, Msg)
```

The same as above, but only on the nodes `Nodes`.

`async_call(Node, Mod, Fun, Args)`

Call `streams` with `promises` is a type of `rpc` which does not suspend the caller until the result is finished. They return a `Key` which can be used at a later stage to collect the value. The key can be viewed as a promise to deliver the answer. The expression `apply(Mod, Fun, Args)` is evaluated for this function on `Node`. Returns `Key` which can be used in a subsequent `yield/1` (see below).

`yield(Key)`

Delivers the promised answer from a previous `async_call` operation. If the answer is available, it is returned immediately. Otherwise, the caller of `yield/1` is suspended until the answer arrives from `Node`.

`nb_yield(Key)`

This is a non-blocking version of `yield`. It returns either the atom `timeout`, or the tuple `{value, V}`, if the computation is finished.

`parallel_eval(ListOfTuples)`

Evaluates the list of size 3 tuples `ListOfTuples`. Each tuple must be of the type `{Mod, Fun, Args}`. Each tuple is sent for evaluation to neighboring nodes, and the replies are collected and returned as a list of individual values. The return values are presented in the same order as the original list `ListOfTuples`.

`pmap({M, F}, Extraargs, List)`

Takes exactly the same arguments and has the same return value as the `lists:map/3` function, except that everything is evaluated in parallel on different nodes.

`pinfo(Pid)`

Location transparent version of `process_info/1`.

`pinfo(Pid, Item)`

Location transparent version of `process_info/2`.

# seq\_trace (Module)

Sequential tracing makes it possible to trace all messages resulting from one initial message. Sequential tracing is completely independent of the ordinary tracing in Erlang, which is controlled by the `erlang:trace/3` BIF. See the chapter "What is Sequential Trace" [page 148] below for more information about what sequential tracing is and how it can be used.

`seq_trace` provides functions which control all aspects of sequential tracing. There are functions for activation, deactivation, inspection and for collection of the trace output.

## Note:

Sequential tracing is currently only available on the jam system and the implementation is in alpha status. This means that the programming interface still might undergo minor adjustments (possibly incompatible) based on feedback from users.

## Exports

```
set_token(Component, ComponentValue) -> {Component, PreviousValue}
```

Types:

- Component = label | serial | Flag
- Flag = send | 'receive' | print | timestamp
- ComponentValue = FlagValue | LabelValue | SerialValue
- FlagValue = bool() (for Flag)
- LabelValue = integer() (for label)
- SerialValue = {Previous, Current}
- Previous = Current = integer()

Sets the individual Component of the trace token to ComponentValue. Returns the previous value of the trace token Component. The valid Component, ComponentValue combinations are:

`label, integer()` The `label` component is an integer which identifies all events belonging to the same sequential trace. If several sequential traces can be active simultaneously `label` is used to identify the separate traces. Default is 0.

`send, bool()` A trace token flag (`true` | `false`) which enables/disables tracing on message sending. Default is `false`.

'receive', bool() A trace token flag (true | false) which enables/disables tracing on message reception. Default is false.

print, bool() A trace token flag (true | false) which enables/disables tracing on explicit calls to seq\_trace:print/1. Default is false.

timestamp, bool() A trace token flag (true | false) which enables/disables a timestamp to be generated for each traced event. Default is false.

serial, SerialValue SerialValue = {Previous, Current}. The serial component contains counters which enables the traced messages to be sorted, should never be set explicitly by the user as these counters are updated automatically. Default is {0, 0}.

set\_token(Token) -> PreviousToken

Types:

- Token = PreviousToken = term() | []

Sets the trace token for the current process to Token. If Token = [] then tracing is disabled, otherwise Token should be an Erlang term returned from get\_token/0 or set\_token/1. set\_token/1 can be used to temporarily exclude message passing from the trace by setting the trace token to empty like this:

```
OldToken = seq_trace:set_token([]), % set to empty and save old value
% do something that should not be part of the trace
io:format("Exclude the signalling caused by this~n"),
seq_trace:set_token(OldToken), % activate the trace token again
...
```

Returns the previous value of the trace token.

get\_token(Component) -> {Component, ComponentValue}

Types:

- Component = label | serial | Flag
- ComponentValue = FlagValue | LabelValue | SerialValue <v>Flag = send | 'receive' | print | timestamp
- FlagValue = bool() (for Flag)
- LabelValue = integer() (for label)
- SerialValue = {Previous, Current} (for serial)
- Previous = Current = integer()

Returns the value of the trace token componentComponent.

get\_token() -> TraceToken

Types:

- TraceToken = term() | []

Returns the value of the trace token for the current process. If [] is returned it means that tracing is not active. Any other value returned is the value of an active trace token. The value returned can be used as input to the set\_token/1 function.

print(TraceInfo) -> void

Types:

- TraceInfo = term()

Puts the Erlang term TraceInfo into the sequential trace output if the process currently is executing within a sequential trace and the print flag of the trace token is set.

`reset_trace()` -> void

Sets the trace token to empty for all processes in the node. The process internal counters used to create the serial of the trace token is set to 0. The trace token is set to empty for all messages in message queues. Together this will effectively stop all ongoing sequential tracing in the Erlang node.

`set_system_tracer(Pid)` -> PreviousPid

Types:

- Pid = PreviousPid = pid() | false

Sets the process denoted by Pid to be the system tracer. Returns the previous value (which can be false if no system tracer is active). The function will generate {'EXIT', {badarg, Info}} if Pid is not the pid of an existing local process.

`get_system_tracer()` -> pid() | false

Returns the pid of the current system tracer or false if no system tracer is activated.

## Trace Messages Sent To the System Tracer

The format of the messages are:

{seq\_trace, Label, SeqTraceInfo, TimeStamp}

or

{seq\_trace, Label, SeqTraceInfo}

depending on whether the timestamp flag of the trace token is set to true or false.

Where :

Label = integer()

TimeStamp = {Seconds, Milliseconds, Microseconds}

Seconds = Milliseconds = Microseconds = integer()

The SeqTraceInfo can have the following formats:

{send, Serial, From, To, Message} Used when a process From with its trace token flag print set to true has sent a message.

{'receive', Serial, From, To, Message} Used when a process To receives a message with a trace token that has the 'receive' flag set to true.

{print, Serial, From, \_, Info} Used when a process From has called `seq_trace:print(Label, Info)` and has a trace token with print set to true and label set to Label.

```
Serial = {PreviousSerial, ThisSerial}
PreviousSerial = ThisSerial = integer()
From = To = pid()
```

`Serial` is a tuple consisting of two integers where the first `PreviousSerial` denotes the serial counter passed in the last received message which carried a trace token. If the process is the first one in a new sequential trace the `PreviousSerial` is set to the value of the process internal “trace clock”. The second integer `ThisSerial` is the serial counter that a process sets on outgoing messages and it is based on the process internal “trace clock” which is incremented by one before it is attached to the trace token in the message.

## What is Sequential Tracing

Sequential tracing is a way to trace a sequence of messages sent between different local or distributed processes where the sequence is initiated by one single message. In short it works like this:

Each process has a *trace token* which can be empty or not empty. When not empty the trace token can be seen as the tuple `{Label, Flags, Serial, From}`. The trace token is passed invisibly with each message.

In order to start a sequential trace the user must explicitly set the trace token in the process that will send the first message in a sequence.

The trace token of a process is automatically set to empty each time the process enters a receive statement but will be set to a value again if the received message carries a nonempty trace token.

On each Erlang node a process can be set as the *system tracer*. This process will receive trace messages each time a message with a trace token is sent or received (if the trace token flag `send` or `'receive'` is set). The system tracer can then print each trace event, write it to a file or whatever suitable.

### Note:

The system tracer will only receive those trace events that occur locally within the Erlang node. To get the whole picture of a sequential trace that involves processes on several Erlang nodes, the output from the system tracer on each involved node must be merged (off line).

In the following sections Sequential Tracing and its most fundamental concepts are described.

## Trace Token

Each process has a current trace token. Initially the token is empty. When the process sends a message to another process, a copy of the current token will be sent “invisibly” along with the message. The current token of a process is set in two ways, either

1. explicitly by the process itself, through a call to `seq_trace:set_token`, or
2. when a message is received.

In both cases the current token will be set. In particular, if the token of a message received is empty, the current token of the process is set to empty.

A trace token contains a label, and a set of flags. Both the label and the flags are set in 1 and 2 above.

## Serial

The trace token contains a component which we call the `Serial` which consists of two integers `Previous` and `Current`. The purpose of `Serial` is uniquely identify each traced event within a trace sequence and to order the messages chronologically and in the different branches if any.

The algorithm for updating `Serial` can be described as follows:

Let each process have two counters `prev_cnt` and `curr_cnt` which both are set to 0 when a process is created. The counters are updated at the following occasions:

- *When the process is about to send a message and the trace token is not empty.*  
Let the `Serial` of the trace token be `tprev` and `tcurr`.  
`curr_cnt := curr_cnt + 1`  
`tprev := prev_cnt`  
`tcurr := curr_cnt`  
The trace token with `tprev` and `tcurr` is then passed along with the message.
- *When the process calls `seq_trace:print(Label, Info)` the `Label` matches the label part of the trace token and the trace token print flag is true.*  
The same algorithm as for send above.
- *When a message is received and contains a nonempty trace token.*  
The process trace token is set to the trace token from the message.  
Let the `Serial` of the trace token be `tprev` and `tcurr`.  
`if (curr_cnt < tcurr)`  
    `curr_cnt := tcurr`  
`prev_cnt := tprev`

The `curr_cnt` of a process is incremented each time the process is involved in a sequential trace. The counter can reach its limit (27 bits) if a process is very long-lived and is involved in much sequential tracing. If the counter overflows it will not be possible to use the `Serial` for ordering of the trace events. To prevent the counter from overflowing in the middle of a sequential trace the function `seq_trace:reset_trace/0` can be called to reset the `prev_cnt` and `curr_cnt` of all processes in the Erlang node. This function will also set all trace tokens in processes and their message queues to empty and will thus stop all ongoing sequential tracing.

## Performance considerations

The performance degradation for a system which is enabled for Sequential tracing is negligible as long as no tracing is activated. When tracing is activated there will of course be an extra cost for each traced message but all other messages will be unaffected.

## Ports

Sequential tracing is not performed across ports.

If the user for some reason wants to pass the trace token to a port this has to be done manually in the code of the port controlling process. The port controlling processes have to check the appropriate sequential trace settings (as obtained from `seq_trace:get_token/1` and include trace information in the message data sent to their respective ports.

Similarly, for messages received from a port, a port controller has to retrieve trace specific information, and set appropriate sequential trace flags through calls to `seq_trace:set_token/2`.

## Distribution

Sequential tracing between nodes is performed transparently. This applies to C-nodes built with `Erl_interface` too. A C-node built with `Erl_interface` only maintains one trace token which means that the C-node will appear as one process from the sequential tracing point of view.

In order to be able to perform sequential tracing between distributed Erlang nodes, the distribution protocol has been extended (in a backward compatible way). An Erlang node which supports sequential tracing can communicate with an older (OTP R3B) node but messages passed within that node can of course not be traced.

## Example of Usage

The example shown here will give rough idea of how the new primitives can be used and what kind of output it will produce.

Assume that we have an initiating process with `Pid = <0.30.0>` like this:

```

-module(seqex).
-compile(export_all).

loop(Port) ->
  receive
    {Port,Message} ->
      seq_trace:set_token(label,17),
      seq_trace:set_token('receive',true),
      seq_trace:set_token(print,true),
      seq_trace:print(17,"**** Trace Started ****"),
      call_server ! {self(),the_message};
    {ack,Ack} ->
      ok
  end,
  loop(Port).

```

And a registered process 'call\_server' with Pid = <0.31.0> like this:

```

loop() ->
  receive
    {PortController,Message} ->
      Ack = {received, Message},
      seq_trace:print(17,"We are here now"),
      PortController ! {ack,Ack}
  end,
  loop().

```

A possible output from the system's sequential\_tracer (inspired by AXE-10 and MD-110) could look like:

```

17:<0.30.0> Info {0,1} WITH
"**** Trace Started ****"
17:<0.31.0> Received {0,2} FROM <0.30.0> WITH
{<0.30.0>,the_message}
17:<0.31.0> Info {2,3} WITH
"We are here now"
17:<0.30.0> Received {2,4} FROM <0.31.0> WITH
{ack,{received,the_message}}

```

The implementation of a system tracer process that produces the printout above could look like this:

```

tracer() ->
  receive
    {seq_trace,Label,TraceInfo} ->
      print_trace(Label,TraceInfo,false);
    {seq_trace,Label,TraceInfo,Ts} ->
      print_trace(Label,TraceInfo,Ts);
    Other -> ignore
  end,
  tracer().

print_trace(Label,TraceInfo,false) ->
  io:format("~p:",[Label]),
  print_trace(TraceInfo);
print_trace(Label,TraceInfo,Ts) ->

```

```

        io:format("~p ~p:", [Label, Ts]),
        print_trace(TraceInfo).

print_trace({print, Serial, From, _, Info}) ->
    io:format("~p Info ~p WITH~n~p~n", [From, Serial, Info]);
print_trace({'receive', Serial, From, To, Message}) ->
    io:format("~p Received ~p FROM ~p WITH~n~p~n", [To, Serial, From, Message]);
print_trace({send, Serial, From, To, Message}) ->
    io:format("~p Sent ~p TO ~p WITH~n~p~n", [From, Serial, To, Message]).

```

The code that creates a process that runs the tracer function above and sets that process as the system tracer could look like this:

```

start() ->
    Pid = spawn(?MODULE, tracer, []),
    seq_trace:set_system_tracer(Pid), % set Pid as the system tracer
    ok.

```

With a function like test/0 below the whole example can be started.

```

test() ->
    P = spawn(?MODULE, loop, [port]),
    register(call_server, spawn(?MODULE, loop, [])),
    start(),
    P ! {port, message}.

```

# user (Module)

`user` is a server which responds to all the messages defined in the I/O interface. The code in `user.er1` can be used as a model for building alternative I/O servers.

## Exports

`start() -> void()`

Starts the basic standard I/O server for the user interface port.

# wrap\_log\_reader (Module)

`wrap_log_reader` is a function to read internally formatted wrap disk logs, refer to `disk_log(3)`. `wrap_log_reader` does not interfere with `disk_log` activities; there is however a known bug in this version of the `wrap_log_reader`, see chapter bugs below.

A wrap disk log file consists of several files, called index files. A log file can be opened and closed. It is also possible to open just one index file separately. If a non-existent or a non-internally formatted file is opened, an error message is returned. If the file is corrupt, no attempt to repair it will be done but an error message is returned.

If a log is configured to be distributed, there is a possibility that all items are not loggen on all nodes. `wrap_log_reader` does only read the log on the called node, it is entirely up to the user to be sure that all items are read.

## Exports

`chunk(Continuation)`

`chunk(Continuation, N) -> {Continuation2, Terms} | {Continuation2, Terms, Badbytes} | {Continuation2, eof} | {error, Reason}`

Types:

- `Continuation = continuation()`
- `N = int() > 0 | infinity`
- `Continuation2 = continuation()`
- `Terms= [term()]`
- `Badbytes = integer()`

This function makes it possible to efficiently read the terms which have been appended to a log. It minimises disk I/O by reading large 8K chunks from the file.

The first time `chunk` is called an initial continuation returned from the `open/1`, `open/2` must be provided.

When `chunk/3` is called, `N` controls the maximum number of terms that are read from the log in each chunk. Default is `infinity`, which means that all the terms contained in the 8K chunk are read. If less than `N` terms are returned, this does not necessarily mean that end of file is reached.

The `chunk` function returns a tuple `{Continuation2, Terms}`, where `Terms` is a list of terms found in the log. `Continuation2` is yet another continuation which must be passed on into any subsequent calls to `chunk`. With a series of calls to `chunk` it is then possible to extract all terms from a log.

The `chunk` function returns a tuple `{Continuation2, Terms, Badbytes}` if the log is opened in read only mode and the read chunk is corrupt. `Badbytes` indicates the number of non-Erlang terms found in the chunk. Note also that the log is not repaired. `chunk` returns `{Continuation2, eof}` when the end of the log is reached, and `{error, Reason}` if an error occurs.

The returned continuation may or may not be valid in the next call to `chunk`. This is because the log may wrap and delete the file into which the continuation points. To make sure this does not happen, the log can be blocked during the search.

```
close(Continuation) -> ok
```

Types:

- `Continuation = continuation()`

This function closes a log file properly.

```
open(Filename) -> OpenRet
```

```
open(Filename, N) -> OpenRet
```

Types:

- `File = string() | atom()`
- `N = integer()`
- `OpenRet = {ok, Continuation} | {error, Reason}`
- `Continuation = continuation()`

`Filename` specifies the name of the file which is to be read.

`N` specifies the index of the file which is to be read. If `N` is omitted the whole wrap log file will be read; if it is specified only the specified index file will be read.

The open function returns `{ok, Continuation}` if the log/index file was successfully opened. The `Continuation` is to be used when chunking or closing the file.

The function returns `{error, Reason}` for all errors.

## Bugs

This version of the `wrap_log_reader` does not detect if the `disk_log` wraps to a new index file between a `wrap_log_reader:open` and the first `wrap_log_reader:chunk`. In this case the chunk will actually read the last logged items in the log file, because the opened index file was truncated by the `disk_log`.

## See Also

`disk_log(3)`

# app (File)

The *application resource file* specifies the resources an application uses, and how the application is started. There must always be one application resource file for each application in the system.

This file is read when an application is loaded, or by the start script generating tools (`systools`).

## FILE SYNTAX

The application resource file is called `Name.app` where `Name` is the name of the application. The file should be located in the `ebin` directory for the application.

The `.app` file contains one single Erlang term, which is called an *application specification*. The file has the following syntax:

```
{application, ApplName,
  [{description, String},
   {vsn,        String},
   {id,         String},
   {modules,    [{Mod1,Vsn1}, Mod2, {Mod3,Vsn3} .., {ModN,VsnN}]},
   {maxP,       Int | infinity},
   {maxT,       Seconds | infinity},
   {registered, [Name1, Name2, ...]},
   {applications, [App1, App2, .., AppN]},
   {included_applications, [App1, App2, .., AppN]},
   {env,        [{Par1, Val1}, {Par2, Val2} .., {ParN, ValN}]},
   {mod,        {Mod, StartArgs}},
   {start_phases, [{Phase, PhaseArgs}]}].
```

The keys have the following meanings:

- `Name = atom()` is the name of the application.
- `Description = string()` is a textual description of the application.
- `Vsn = string()` is the version of the application. This string must be a valid filename.
- `Id = string()` is the product identification of the application.
- `Modules = [Mod1 | {Mod1, Vsn1}]` is a list of all the modules and their versions introduced by this application. A module can be listed without version, only the name of the module is stated. A module can only be defined in one application.
- `MaxT = int() | infinity` is the maximum time that the application can run (or the atom `infinity`). The key `maxT` is optional and defaults to `infinity`.

- `Registered = [atom()]` is a list of all the names of registered processes started in this application.
- `applications = [atom()]` is a list of applications which must be started before this application is started. Most applications have dependencies to the `kernel` and `stdlib` applications.
- `included_applications = [atom()]` is a list of applications which are included by this application. An included application is loaded, but not started, by the `application_controller`. Processes implemented in an included application should be placed underneath a supervisor in the including application. This key is optional and defaults to `[]`.
- `env` is a list of the environment variables in the application. Each parameter `ParX` is an atom, and the associated `ValX` can be any term. The `env` key is optional and defaults to an empty list.
- `mod` is the application callback module of the application behaviour. The application master starts the application by evaluating the function `Mod:start(Type, StartArgs)`. When the application has stopped, by command or because it terminates, the application master calls `Mod:stop(State)` to let the application clean up. If no `State` was returned from `Mod:start/2`, `Mod:stop([])` is called.  
The `mod` key should be omitted for applications which are code libraries, such as the `stdlib` application. These applications have no dynamic behaviour of their own and should not have a start function.
- `start_phases` is a list of start phases and the attached start arguments for the application. The application master starts the application by evaluating the function `Mod:start_phase(Phase, Type, PhaseArgs)` for each defined start phase. `Mod` is the same callback module as defined in the `mod` key. Each parameter `Phase` is an atom, and the associated `PhaseArgs` is a list of any terms. The key `start_phases` is optional, and the behaviour of the system is dependent if the key is defined or not, refer to application (3).

## SEE ALSO

application(3), systools(3)

# config (File)

A *configuration file* contains values for configuration parameters for the applications in the system. The command `erl -config Name` tells the system to use data in the system configuration file `Name.config` to override the arguments contained in the application resource files for the set of applications used by this system.

An application should call `application:get_env(AppName, Parameter)` to retrieve the values for the configuration parameters.

The parameters can also be overridden from the command line:

```
erl -AppName Par1 Val1 Par2 Val2 ...
```

## Note:

Each term should be an Erlang term. However, in the Unix shell, the term must be enclosed in single quotation marks. For example: `{file, "a.log"}`.

## FILE SYNTAX

The configuration file is called `Name.config` where `Name` is the name of the application.

The file has the following syntax:

```
[{AppName, [{Par, Val}]}].
```

There is one tuple for each application. The second element in each tuple is a list of configuration parameters and their values.

- `AppName = atom()` is the name of the application.
- `Par = atom()` is the name of a configuration parameter.
- `Val = term()` is the value of the configuration parameter.

## SEE ALSO

`application(3)`, `systools(3)`

# Index

Modules are typed in *this* way.  
Functions are typed in *this* way.

- abcast/2
  - rpc* , 142
- abcast/3
  - rpc* , 142
- abs/1
  - erlang* , 63
- accept/1
  - gen\_tcp* , 108
- accept/2
  - gen\_tcp* , 108
- add\_path/1
  - code* , 39
- add\_patha/1
  - code* , 39
- add\_paths/1
  - code* , 39
- add\_pathsa/1
  - code* , 40
- add\_pathsz/1
  - code* , 39
- add\_pathz/1
  - code* , 39
- add\_report\_handler/1
  - error\_logger* , 92
- add\_report\_handler/2
  - error\_logger* , 92
- add\_slave/1
  - erl\_boot\_server* , 54
- all\_loaded/0
  - code* , 42
- allow/1
  - net\_kernel* , 136
- alog/2
  - disk\_log* , 46
- alog\_terms/2
  - disk\_log* , 47
- application*
  - get\_all\_env*/0, 28
  - get\_all\_env*/1, 28
  - get\_all\_key*/0, 28
  - get\_all\_key*/1, 28
  - get\_application*/0, 29
  - get\_application*/1, 29
  - get\_env*/1, 29
  - get\_env*/2, 29
  - get\_key*/1, 29
  - get\_key*/2, 29
  - load*/1, 29
  - load*/2, 29
  - loaded\_applications*/0, 30
  - Module:config\_change*/3, 32
  - Module:start*/2, 33
  - Module:start\_phase*/3, 33
  - Module:stop*/1, 34
  - permit*/2, 30
  - start*/1, 31
  - start*/2, 31
  - start\_type*/0, 31
  - stop*/1, 31
  - takeover*/2, 32
  - which\_applications*/0, 32
- apply/3
  - erlang* , 63
- async\_call/4
  - rpc* , 144
- atom\_to\_list/1
  - erlang* , 63
- auth*
  - cookie*/0, 37
  - cookie*/1, 37
  - exists*/1, 36

- is\_auth/1, 36
- node\_cookie/2, 37
- open/1, 36
- start/0, 36
- stop/0, 36
  
- balog/2
  - disk\_log*, 46
- balog\_terms/2
  - disk\_log*, 47
- binary\_to\_list/1
  - erlang*, 63
- binary\_to\_list/3
  - erlang*, 63
- binary\_to\_term/1
  - erlang*, 64
- block/1
  - disk\_log*, 47
- block/2
  - disk\_log*, 47
- block\_call/4
  - rpc*, 141
- blog/2
  - disk\_log*, 49
- blog\_terms/2
  - disk\_log*, 50
- boot/1
  - init*, 130
- breopen/3
  - disk\_log*, 52
- btruncate/2
  - disk\_log*, 52
  
- call/4
  - rpc*, 141
- cast/4
  - rpc*, 141
- change\_group/2
  - file*, 100
- change\_owner/2
  - file*, 100
- change\_owner/3
  - file*, 100
- change\_size/2
  - disk\_log*, 47
- change\_time/2
  - file*, 100
- change\_time/3
  - file*, 100
- chunk/1
  - wrap\_log\_reader*, 154
- chunk/2
  - disk\_log*, 48
  - wrap\_log\_reader*, 154
- chunk/3
  - disk\_log*, 48
- chunk\_step/3
  - disk\_log*, 49
- clash/0
  - code*, 44
- clear\_cmd/0
  - heart*, 122
- close/1
  - disk\_log*, 49
  - file*, 103
  - gen\_tcp*, 108
  - gen\_udp*, 111
  - inet*, 124
  - wrap\_log\_reader*, 155
- cmd/1
  - os*, 137
- code
  - add\_path/1, 39
  - add\_patha/1, 39
  - add\_paths/1, 39
  - add\_pathsa/1, 40
  - add\_pathsz/1, 39
  - add\_pathz/1, 39
  - all\_loaded/0, 42
  - clash/0, 44
  - compiler\_dir/0, 43
  - del\_path/1, 40
  - delete/1, 41
  - delete\_interpret/1, 44
  - ensure\_loaded/1, 41
  - get\_object\_code/1, 43
  - get\_path/0, 39
  - interpret/1, 44
  - interpret\_binary/3, 44
  - interpreted/0, 45
  - interpreted/1, 45
  - is\_loaded/1, 41

- lib\_dir/0, 42
- lib\_dir/1, 42
- load\_abs/1, 40
- load\_binary/3, 42
- load\_file/1, 40
- objfile\_extension/0, 43
- priv\_dir/1, 43
- purge/1, 41
- replace\_path/2, 40
- root\_dir/0, 42
- set\_path/1, 39
- soft\_purge/1, 41
- start/0, 38
- start/1, 38
- start\_link/0, 38
- start\_link/1, 38
- stick\_dir/1, 43
- stop/0, 42
- unstick\_dir/1, 44
- which/1, 44
- compiler\_dir/0
  - code, 43
- concat\_binary/1
  - erlang, 64
- connect/3
  - gen\_tcp, 108
- connect/4
  - gen\_tcp, 108
- consult/1
  - file, 101
- controlling\_process/2
  - gen\_tcp, 109
  - gen\_udp, 111
- cookie/0
  - auth, 37
- cookie/1
  - auth, 37
- create/1
  - pg2, 139
- date/0
  - erlang, 64
- del\_dir/1
  - file, 96
- del\_lock/1
  - global, 113
- del\_lock/2
  - global, 113
- del\_path/1
  - code, 40
- delete/1
  - code, 41
  - file, 95
  - pg2, 139
- delete\_interpret/1
  - code, 44
- delete\_report\_handler/1
  - error\_logger, 92
- delete\_slave/1
  - erl\_boot\_server, 55
- disk\_log
  - alog/2, 46
  - alog\_terms/2, 47
  - balog/2, 46
  - balog\_terms/2, 47
  - block/1, 47
  - block/2, 47
  - blog/2, 49
  - blog\_terms/2, 50
  - breopen/3, 52
  - btruncate/2, 52
  - change\_size/2, 47
  - chunk/2, 48
  - chunk/3, 48
  - chunk\_step/3, 49
  - close/1, 49
  - inc\_wrap\_log/1, 49
  - log/2, 49
  - log\_terms/2, 50
  - open/1, 50
  - reopen/2, 52
  - reopen/3, 52
  - sync/1, 52
  - truncate/1, 52
  - truncate/2, 52
  - unblock/1, 52
- dns\_hostname/1
  - net\_adm, 134
- element/2
  - erlang, 65
- ensure\_loaded/1
  - code, 41
- erase/0
  - erlang, 65

- erase/1
  - erlang, 65
- erl\_boot\_server
  - add\_slave/1, 54
  - delete\_slave/1, 55
  - start/1, 54
  - start\_link/1, 54
  - which\_slaves/0, 55
- erl\_dll
  - format\_error/1, 57
  - load\_driver/2, 56
  - loaded\_drivers/0, 57
  - start/0, 56
  - start\_link/0, 56
  - stop/0, 56
  - unload\_driver/1, 56
- erl\_prim\_loader
  - get\_file/1, 59
  - get\_path/0, 60
  - set\_path/1, 60
  - start/3, 59
- erlang
  - abs/1, 63
  - apply/3, 63
  - atom\_to\_list/1, 63
  - binary\_to\_list/1, 63
  - binary\_to\_list/3, 63
  - binary\_to\_term/1, 64
  - concat\_binary/1, 64
  - date/0, 64
  - element/2, 65
  - erase/0, 65
  - erase/1, 65
  - erlang:cancel\_timer/1, 64
  - erlang:check\_process\_code/2, 64
  - erlang:delete\_module/1, 64
  - erlang:demonitor/1, 65
  - erlang:disconnect\_node/1, 65
  - erlang:get\_cookie/0, 67
  - erlang:hash/2, 68
  - erlang:info/1, 68
  - erlang:load\_module/2, 70
  - erlang:localtime/0, 71
  - erlang:localtime\_to\_universaltime/1, 71
  - erlang:module\_loaded/1, 71
  - erlang:monitor/2, 72
  - erlang:port\_info/2, 75
  - erlang:ports/0, 75
  - erlang:pre\_loaded/0, 76
  - erlang:process\_display/2, 76
  - erlang:purge\_module/1, 78
  - erlang:send\_after/3, 79
  - erlang:set\_cookie/2, 79
  - erlang:start\_timer/3, 81
  - erlang:system\_flag/2, 82
  - erlang:universaltime/0, 85
  - erlang:universaltime\_to\_localtime/1, 85
  - exit/1, 66
  - exit/2, 66
  - float/1, 66
  - float\_to\_list/1, 67
  - get/0, 67
  - get/1, 67
  - get\_keys/1, 67
  - group\_leader/0, 68
  - group\_leader/2, 68
  - halt/0, 68
  - halt/1, 68
  - hd/1, 68
  - integer\_to\_list/1, 69
  - is\_alive/0, 69
  - is\_process\_alive/1, 69
  - length/1, 69
  - link/1, 69
  - list\_to\_atom/1, 69
  - list\_to\_binary/1, 69
  - list\_to\_float/1, 70
  - list\_to\_integer/1, 70
  - list\_to\_pid/1, 70
  - list\_to\_tuple/1, 70
  - make\_ref/0, 71
  - monitor\_node/2, 72
  - node/0, 72
  - node/1, 73
  - nodes/0, 73
  - now/0, 73
  - open\_port/2, 73
  - pid\_to\_list/1, 75
  - process\_flag/2, 76
  - process\_info/1, 76
  - process\_info/2, 77
  - processes/0, 77
  - put/2, 78
  - register/2, 78
  - registered/0, 78
  - round/1, 79
  - self/0, 79
  - setelement/3, 79
  - size/1, 79
  - spawn/3, 80
  - spawn/4, 80
  - spawn\_link/3, 80

- spawn\_link/4, 80
- spawn\_opt/4, 80
- split\_binary/2, 81
- statistics/1, 81
- term\_to\_binary/1, 82
- throw/1, 82
- time/0, 83
- tl/1, 83
- trace/3, 83
- trunc/1, 84
- tuple\_to\_list/1, 85
- unlink/1, 85
- unregister/1, 85
- whereis/1, 86
- erlang:cancel\_timer/1  
*erlang*, 64
- erlang:check\_process\_code/2  
*erlang*, 64
- erlang:delete\_module/1  
*erlang*, 64
- erlang:demonitor/1  
*erlang*, 65
- erlang:disconnect\_node/1  
*erlang*, 65
- erlang:get\_cookie/0  
*erlang*, 67
- erlang:hash/2  
*erlang*, 68
- erlang:info/1  
*erlang*, 68
- erlang:load\_module/2  
*erlang*, 70
- erlang:localtime/0  
*erlang*, 71
- erlang:localtime\_to\_universaltime/1  
*erlang*, 71
- erlang:module\_loaded/1  
*erlang*, 71
- erlang:monitor/2  
*erlang*, 72
- erlang:port\_info/2  
*erlang*, 75
- erlang:ports/0  
*erlang*, 75
- erlang:pre\_loaded/0
- erlang*, 76
- erlang:process\_display/2  
*erlang*, 76
- erlang:purge\_module/1  
*erlang*, 78
- erlang:send\_after/3  
*erlang*, 79
- erlang:set\_cookie/2  
*erlang*, 79
- erlang:start\_timer/3  
*erlang*, 81
- erlang:system\_flag/2  
*erlang*, 82
- erlang:universaltime/0  
*erlang*, 85
- erlang:universaltime\_to\_localtime/1  
*erlang*, 85
- error\_handler*
  - undefined\_function/3, 87
  - undefined\_lambda/3, 87
- error\_logger*
  - add\_report\_handler/1, 92
  - add\_report\_handler/2, 92
  - delete\_report\_handler/1, 92
  - error\_msg/1, 91
  - error\_msg/2, 91
  - error\_report/1, 89
  - error\_report/2, 90
  - format/2, 91
  - info\_msg/1, 91
  - info\_msg/2, 91
  - info\_report/1, 90
  - info\_report/2, 90
  - logfile/1, 91
  - start/0, 89
  - start\_link/0, 89
  - swap\_handler/1, 92
  - tty/1, 91
- error\_msg/1  
*error\_logger*, 91
- error\_msg/2  
*error\_logger*, 91
- error\_report/1  
*error\_logger*, 89
- error\_report/2  
*error\_logger*, 90

- eval/1
  - file*, 101
- eval\_everywhere/3
  - rpc*, 142
- eval\_everywhere/4
  - rpc*, 142
- exists/1
  - auth*, 36
- exit/1
  - erlang*, 66
- exit/2
  - erlang*, 66
- file*
  - change\_group*/2, 100
  - change\_owner*/2, 100
  - change\_owner*/3, 100
  - change\_time*/2, 100
  - change\_time*/3, 100
  - close*/1, 103
  - consult*/1, 101
  - del\_dir*/1, 96
  - delete*/1, 95
  - eval*/1, 101
  - file\_info*/1, 99
  - format\_error*/1, 104
  - get\_cwd*/0, 95
  - get\_cwd*/1, 95
  - list\_dir*/1, 97
  - make\_dir*/1, 96
  - make\_link*/2, 97
  - make\_symlink*/2, 97
  - open*/2, 102
  - path\_consult*/2, 101
  - path\_eval*/2, 101
  - path\_open*/3, 103
  - position*/2, 104
  - pread*/3, 104
  - pwrite*/3, 104
  - read*/2, 103
  - read\_file*/1, 94
  - read\_file\_info*/1, 97
  - read\_link*/1, 99
  - read\_link\_info*/1, 98
  - rename*/2, 96
  - set\_cwd*/1, 95
  - sync*/1, 104
  - truncate*/1, 104
  - write*/2, 103
  - write\_file*/2, 94
  - write\_file\_info*/2, 99
- file\_info*/1
  - file*, 99
- find\_executable*/1
  - os*, 137
- find\_executable*/2
  - os*, 137
- float*/1
  - erlang*, 66
- float\_to\_list*/1
  - erlang*, 67
- format*/2
  - error\_logger*, 91
- format\_error*/1
  - erl\_dll*, 57
  - file*, 104
  - inet*, 123
- gen\_tcp*
  - accept*/1, 108
  - accept*/2, 108
  - close*/1, 108
  - connect*/3, 108
  - connect*/4, 108
  - controlling\_process*/2, 109
  - listen*/2, 109
  - recv*/2, 109
  - recv*/3, 109
  - send*/2, 110
- gen\_udp*
  - close*/1, 111
  - controlling\_process*/2, 111
  - open*/1, 111
  - open*/2, 111
  - recv*/2, 112
  - recv*/3, 112
  - send*/4, 112
- get*/0
  - erlang*, 67
- get*/1
  - erlang*, 67
- get\_all\_env*/0
  - application*, 28
- get\_all\_env*/1
  - application*, 28
- get\_all\_key*/0

- application* , 28
- get\_all\_key/1
  - application* , 28
- get\_application/0
  - application* , 29
- get\_application/1
  - application* , 29
- get\_args/0
  - init* , 131
- get\_argument/1
  - init* , 131
- get\_arguments/0
  - init* , 130
- get\_closest\_pid/1
  - pg2* , 139
- get\_cwd/0
  - file* , 95
- get\_cwd/1
  - file* , 95
- get\_env/1
  - application* , 29
- get\_env/2
  - application* , 29
- get\_file/1
  - erl\_prim\_loader* , 59
- get\_key/1
  - application* , 29
- get\_key/2
  - application* , 29
- get\_keys/1
  - erlang* , 67
- get\_local\_members/1
  - pg2* , 140
- get\_members/1
  - pg2* , 139
- get\_object\_code/1
  - code* , 43
- get\_path/0
  - code* , 39
  - erl\_prim\_loader* , 60
- get\_status/0
  - init* , 131
- get\_system\_tracer/0
  - seq\_trace* , 147
- get\_token/0
  - seq\_trace* , 146
- get\_token/1
  - seq\_trace* , 146
- getaddr/2
  - inet* , 125
- getenv/1
  - os* , 137
- gethostbyaddr/1
  - inet* , 123
- gethostbyname/1
  - inet* , 123
- gethostbyname/2
  - inet* , 124
- gethostname/0
  - inet* , 124
- getpid/0
  - os* , 138
- global*
  - del\_lock*/1, 113
  - del\_lock*/2, 113
  - notify\_all\_name*/3, 114
  - random\_exit\_name*/3, 114
  - random\_notify\_name*/3, 114
  - re\_register\_name*/2, 115
  - re\_register\_name*/3, 115
  - register\_name*/2, 114
  - register\_name*/3, 114
  - registered\_names*/0, 114
  - send*/2, 115
  - set\_lock*/1, 115
  - set\_lock*/2, 115
  - set\_lock*/3, 115
  - start*/0, 116
  - start\_link*/0, 116
  - stop*/0, 116
  - sync*/0, 116
  - trans*/2, 116
  - trans*/3, 116
  - trans*/4, 116
  - unregister\_name*/1, 116
  - whereis\_name*/1, 116
- global\_group*
  - global\_groups*/0, 118
  - info*/0, 118
  - monitor\_nodes*/1, 118
  - own\_nodes*/0, 118

registered\_names/2, 118  
send/2, 119  
send/4, 119  
start/0, 120  
start\_link/0, 120  
stop/0, 120  
sync/0, 119  
whereis\_name/1, 119  
whereis\_name/3, 119

global\_groups/0  
  *global\_group*, 118

group\_leader/0  
  *erlang*, 68

group\_leader/2  
  *erlang*, 68

halt/0  
  *erlang*, 68

halt/1  
  *erlang*, 68

hd/1  
  *erlang*, 68

*heart*  
  clear\_cmd/0, 122  
  set\_cmd/1, 122  
  start/0, 121

host\_file/0  
  *net\_adm*, 134

inc\_wrap\_log/1  
  *disk\_log*, 49

*inet*  
  close/1, 124  
  format\_error/1, 123  
  getaddr/2, 125  
  gethostbyaddr/1, 123  
  gethostbyname/1, 123  
  gethostbyname/2, 124  
  gethostname/0, 124  
  peername/1, 124  
  port/1, 124  
  setopts/2, 125  
  sockname/1, 124

info/0  
  *global\_group*, 118

info\_msg/1  
  *error\_logger*, 91

info\_msg/2  
  *error\_logger*, 91

info\_report/1  
  *error\_logger*, 90

info\_report/2  
  *error\_logger*, 90

*init*  
  boot/1, 130  
  get\_args/0, 131  
  get\_argument/1, 131  
  get\_arguments/0, 130  
  get\_status/0, 131  
  reboot/0, 131  
  restart/0, 131  
  script\_id/0, 132  
  stop/0, 131

integer\_to\_list/1  
  *erlang*, 69

interpret/1  
  *code*, 44

interpret\_binary/3  
  *code*, 44

interpreted/0  
  *code*, 45

interpreted/1  
  *code*, 45

is\_alive/0  
  *erlang*, 69

is\_auth/1  
  *auth*, 36

is\_loaded/1  
  *code*, 41

is\_process\_alive/1  
  *erlang*, 69

join/2  
  *pg2*, 140

kernel\_apply/3  
  *net\_kernel*, 136

leave/2  
  *pg2*, 140

length/1  
  *erlang*, 69

lib\_dir/0

- code*, 42
- lib\_dir/1
  - code*, 42
- link/1
  - erlang*, 69
- list\_dir/1
  - file*, 97
- list\_to\_atom/1
  - erlang*, 69
- list\_to\_binary/1
  - erlang*, 69
- list\_to\_float/1
  - erlang*, 70
- list\_to\_integer/1
  - erlang*, 70
- list\_to\_pid/1
  - erlang*, 70
- list\_to\_tuple/1
  - erlang*, 70
- listen/2
  - gen\_tcp*, 109
- load/1
  - application*, 29
- load/2
  - application*, 29
- load\_abs/1
  - code*, 40
- load\_binary/3
  - code*, 42
- load\_driver/2
  - erl\_ddll*, 56
- load\_file/1
  - code*, 40
- loaded\_applications/0
  - application*, 30
- loaded\_drivers/0
  - erl\_ddll*, 57
- localhost/0
  - net\_adm*, 134
- log/2
  - disk\_log*, 49
- log\_terms/2
  - disk\_log*, 50
- logfile/1
  - error\_logger*, 91
- make\_dir/1
  - file*, 96
- make\_link/2
  - file*, 97
- make\_ref/0
  - erlang*, 71
- make\_symlink/2
  - file*, 97
- Module:config\_change/3
  - application*, 32
- Module:start/2
  - application*, 33
- Module:start\_phase/3
  - application*, 33
- Module:stop/1
  - application*, 34
- monitor\_node/2
  - erlang*, 72
- monitor\_nodes/1
  - global\_group*, 118
  - net\_kernel*, 136
- multi\_server\_call/2
  - rpc*, 143
- multi\_server\_call/3
  - rpc*, 143
- multicall/3
  - rpc*, 142
- multicall/4
  - rpc*, 143
- names/0
  - net\_adm*, 134
- nb\_yield/1
  - rpc*, 144
- net\_adm*
  - dns\_hostname*/1, 134
  - host\_file*/0, 134
  - localhost*/0, 134
  - names*/0, 134
  - ping*/1, 134
  - world*/0, 134
  - world\_list*/2, 134

*net\_kernel*  
  allow/1, 136  
  kernel\_apply/3, 136  
  monitor\_nodes/1, 136

node/0  
  erlang, 72

node/1  
  erlang, 73

node\_cookie/2  
  auth, 37

nodes/0  
  erlang, 73

notify\_all\_name/3  
  global, 114

now/0  
  erlang, 73

objfile\_extension/0  
  code, 43

open/1  
  auth, 36  
  disk\_log, 50  
  gen\_udp, 111  
  wrap\_log\_reader, 155

open/2  
  file, 102  
  gen\_udp, 111  
  wrap\_log\_reader, 155

open\_port/2  
  erlang, 73

os  
  cmd/1, 137  
  find\_executable/1, 137  
  find\_executable/2, 137  
  getenv/1, 137  
  getpid/0, 138  
  type/0, 138  
  version/0, 138

own\_nodes/0  
  global\_group, 118

parallel\_eval/1  
  rpc, 144

path\_consult/2  
  file, 101

path\_eval/2  
  file, 101

path\_open/3  
  file, 103

peername/1  
  inet, 124

permit/2  
  application, 30

pg2  
  create/1, 139  
  delete/1, 139  
  get\_closest\_pid/1, 139  
  get\_local\_members/1, 140  
  get\_members/1, 139  
  join/2, 140  
  leave/2, 140  
  start/0, 140  
  start\_link/0, 140  
  which\_groups/0, 140

pid\_to\_list/1  
  erlang, 75

pinfo/1  
  rpc, 144

pinfo/2  
  rpc, 144

ping/1  
  net\_adm, 134

pmap/4  
  rpc, 144

port/1  
  inet, 124

position/2  
  file, 104

pread/3  
  file, 104

print/1  
  seq\_trace, 146

priv\_dir/1  
  code, 43

process\_flag/2  
  erlang, 76

process\_info/1  
  erlang, 76

process\_info/2  
  erlang, 77

---

processes/0  
     *erlang*, 77

purge/1  
     *code*, 41

put/2  
     *erlang*, 78

pwrite/3  
     *file*, 104

random\_exit\_name/3  
     *global*, 114

random\_notify\_name/3  
     *global*, 114

re\_register\_name/2  
     *global*, 115

re\_register\_name/3  
     *global*, 115

read/2  
     *file*, 103

read\_file/1  
     *file*, 94

read\_file\_info/1  
     *file*, 97

read\_link/1  
     *file*, 99

read\_link\_info/1  
     *file*, 98

reboot/0  
     *init*, 131

recv/2  
     *gen\_tcp*, 109  
     *gen\_udp*, 112

recv/3  
     *gen\_tcp*, 109  
     *gen\_udp*, 112

register/2  
     *erlang*, 78

register\_name/2  
     *global*, 114

register\_name/3  
     *global*, 114

registered/0  
     *erlang*, 78

registered\_names/0  
     *global*, 114

registered\_names/2  
     *global\_group*, 118

rename/2  
     *file*, 96

reopen/2  
     *disk\_log*, 52

reopen/3  
     *disk\_log*, 52

replace\_path/2  
     *code*, 40

reset\_trace/0  
     *seq\_trace*, 147

restart/0  
     *init*, 131

root\_dir/0  
     *code*, 42

round/1  
     *erlang*, 79

rpc
 

- abcast/2, 142
- abcast/3, 142
- async\_call/4, 144
- block\_call/4, 141
- call/4, 141
- cast/4, 141
- eval\_everywhere/3, 142
- eval\_everywhere/4, 142
- multi\_server\_call/2, 143
- multi\_server\_call/3, 143
- multicall/3, 142
- multicall/4, 143
- nb\_yield/1, 144
- parallel\_eval/1, 144
- pinfo/1, 144
- pinfo/2, 144
- pmap/4, 144
- safe\_multi\_server\_call/2, 143
- safe\_multi\_server\_call/3, 143
- sbcast/2, 142
- sbcast/3, 142
- server\_call/4, 142
- start/0, 141
- stop/0, 141
- yield/1, 144

safe\_multi\_server\_call/2  
     *rpc*, 143

- safe\_multi\_server\_call/3
  - rpc*, 143
- sbcast/2
  - rpc*, 142
- sbcast/3
  - rpc*, 142
- script\_id/0
  - init*, 132
- self/0
  - erlang*, 79
- send/2
  - gen\_tcp*, 110
  - global*, 115
  - global\_group*, 119
- send/4
  - gen\_udp*, 112
  - global\_group*, 119
- seq\_trace
  - get\_system\_tracer*/0, 147
  - get\_token*/0, 146
  - get\_token*/1, 146
  - print*/1, 146
  - reset\_trace*/0, 147
  - set\_system\_tracer*/1, 147
  - set\_token*/1, 146
  - set\_token*/2, 145
- server\_call/4
  - rpc*, 142
- set\_cmd/1
  - heart*, 122
- set\_cwd/1
  - file*, 95
- set\_lock/1
  - global*, 115
- set\_lock/2
  - global*, 115
- set\_lock/3
  - global*, 115
- set\_path/1
  - code*, 39
  - erl\_prim\_loader*, 60
- set\_system\_tracer/1
  - seq\_trace*, 147
- set\_token/1
  - seq\_trace*, 146
- set\_token/2
  - seq\_trace*, 145
- setelement/3
  - erlang*, 79
- setopts/2
  - inet*, 125
- size/1
  - erlang*, 79
- sockname/1
  - inet*, 124
- soft\_purge/1
  - code*, 41
- spawn/3
  - erlang*, 80
- spawn/4
  - erlang*, 80
- spawn\_link/3
  - erlang*, 80
- spawn\_link/4
  - erlang*, 80
- spawn\_opt/4
  - erlang*, 80
- split\_binary/2
  - erlang*, 81
- start/0
  - auth*, 36
  - code*, 38
  - erl\_ddll*, 56
  - error\_logger*, 89
  - global*, 116
  - global\_group*, 120
  - heart*, 121
  - pg2*, 140
  - rpc*, 141
  - user*, 153
- start/1
  - application*, 31
  - code*, 38
  - erl\_boot\_server*, 54
- start/2
  - application*, 31
- start/3
  - erl\_prim\_loader*, 59
- start\_link/0
  - code*, 38

---

- erl\_dll* , 56
- error\_logger* , 89
- global* , 116
- global\_group* , 120
- pg2* , 140
- start\_link/1*
  - code* , 38
  - erl\_boot\_server* , 54
- start\_type/0*
  - application* , 31
- statistics/1*
  - erlang* , 81
- stick\_dir/1*
  - code* , 43
- stop/0*
  - auth* , 36
  - code* , 42
  - erl\_dll* , 56
  - global* , 116
  - global\_group* , 120
  - init* , 131
  - rpc* , 141
- stop/1*
  - application* , 31
- swap\_handler/1*
  - error\_logger* , 92
- sync/0*
  - global* , 116
  - global\_group* , 119
- sync/1*
  - disk\_log* , 52
  - file* , 104
- takeover/2*
  - application* , 32
- term\_to\_binary/1*
  - erlang* , 82
- throw/1*
  - erlang* , 82
- time/0*
  - erlang* , 83
- tl/1*
  - erlang* , 83
- trace/3*
  - erlang* , 83
- trans/2*
  - global* , 116
- trans/3*
  - global* , 116
- trans/4*
  - global* , 116
- trunc/1*
  - erlang* , 84
- truncate/1*
  - disk\_log* , 52
  - file* , 104
- truncate/2*
  - disk\_log* , 52
- tty/1*
  - error\_logger* , 91
- tuple\_to\_list/1*
  - erlang* , 85
- type/0*
  - os* , 138
- unblock/1*
  - disk\_log* , 52
- undefined\_function/3*
  - error\_handler* , 87
- undefined\_lambda/3*
  - error\_handler* , 87
- unlink/1*
  - erlang* , 85
- unload\_driver/1*
  - erl\_dll* , 56
- unregister/1*
  - erlang* , 85
- unregister\_name/1*
  - global* , 116
- unstick\_dir/1*
  - code* , 44
- user*
  - start/0* , 153
- version/0*
  - os* , 138
- whereis/1*
  - erlang* , 86
- whereis\_name/1*

- global* , 116
  - global\_group* , 119
- whereis\_name/3
  - global\_group* , 119
- which/1
  - code , 44
- which\_applications/0
  - application , 32
- which\_groups/0
  - pg2 , 140
- which\_slaves/0
  - erl\_boot\_server , 55
- world/0
  - net\_adm , 134
- world\_list/2
  - net\_adm , 134
- wrap\_log\_reader
  - chunk/1, 154
  - chunk/2, 154
  - close/1, 155
  - open/1, 155
  - open/2, 155
- write/2
  - file , 103
- write\_file/2
  - file , 94
- write\_file\_info/2
  - file , 99
- yield/1
  - rpc , 144