

Getting Started

version 4.8

OTP Team

1998-01-22

Typeset in \LaTeX from SGML source using the DOCBUILDER 3.0 Document System.

Contents

1	Getting Started	1
1.1	Getting Started with Erlang	2
	Introduction	2
	Working with Erlang	3
	The Erlang Shell	6
	Compiling	14
	Debugging with the Process Manager	16
	Distributed Erlang	18
	Configuration	22
	List of Figures	23
	List of Tables	25

Chapter 1

Getting Started

1.1 Getting Started with Erlang

Introduction

This chapter describes the most common tasks users perform in the *Erlang Development Environment* or *Open Telecom Platform* and is intended for readers who are running Erlang for the first time. It also contains references to more detailed information.

This manual assumes that you have another source of information about how the language Erlang works. The first chapter of the book [erlbook2] is available online in PDF¹ format.

This manual will describe how to

- compile and run Erlang programs
- work with the *Erlang shell*, this is the most central part of the development environment
- use the basic tools, the debugger and process manager

This manual will also describe differences in development environment on the Unix and Windows platforms.

Note:

This manual will not describe how to use *xerl*. Xerl is a more integrated development environment, with an integrated editor and an Erlang shell. Xerl has some buttons and menus to control the Erlang shell, to start other tools and to find manual pages. This is an old tool that only work on Unix systems and is currently included for backward compatibility with older releases. It will be removed or completely rewritten in a future release.

Manual conventions

The following conventions are used to describe commands, which are entered from the keyboard in the Erlang shell, or on the Emacs command line: window-based

- to enter the command *C-c*: Press the *Control key* and the letter *c* simultaneously. *C-c* is equivalent to $\sim c$.
- to enter the command *M-f*: Press and release the *Esc key*, then press the letter *f*.
- to enter the command *halt()*. Enter *halt()*. and then press the *Return key*. You must include the full stop that terminates the command.
- the Unix prompt is *>*.
- the dos prompt is *C:\>*

¹URL: <http://www.erlang.org/download/erlang-book-part1.pdf>

Working with Erlang

Erlang programs are executed when you instruct the *Erlang Runtime System* to execute your code. They do not run as binary programs, which can be executed directly by your computer. The Erlang Runtime System consists of an Erlang evaluator and Erlang libraries. The Erlang evaluator is often referred to as an *emulator* and is very similar to the Java virtual machine.

The Erlang Runtime system together with a basic set of tools for software development form the product *The Erlang Development Environment*. On top of the Erlang Development Environment, an advanced set of tools can be added to form the *The Open Telecom Platform*.

The most central part of the Erlang development environment is the *Erlang shell*, *erl*. *erl* is available on Unix systems, Windows NT and Windows 95. It looks very similar to a Unix shell, Windows DOS box or a Windows NT Command box. The difference is that the Erlang command line interface understands how to

- compile and load Erlang programs
- run Erlang programs, individual functions or evaluate Erlang expressions
- monitor and control execution of the programs
- command line editing
- history list with previously entered commands

In addition to this text-based interface there are several window-based *tools* like debugger, process manager and database visualizer.

You can access everything from this command interface but you can also activate the graphic *toolbar* where you get buttons to start the window based tools.

Running Erlang

To start the Erlang development environment on a Unix system you execute in a Unix shell

```
> erl -s toolbar
```

On the Windows platform you should find Erlang in the start menu, use the entry Erlang with Toolbar. If the Erlang executable is in your execution path you can also start Erlang from a DOS box or NT command line with

```
C:\> werl -s toolbar
```

Note:

On the Windows platform there are two executables, *erl* and *werl*. You normally use *werl*, see the *werl* reference manual for details.

The following appears on the screen:

```
Erlang (JAM) emulator version 4.6
```

```
Eshell V4.6 (abort with ^G)
1>
```

The Erlang shell is now ready for you to type some input. The commands you enter could be any Erlang expression, which is understood by an Erlang shell.

Try it out by typing `1 + 2.` followed by pressing the return key (don't forget the terminating dot). You can also call one of the built in functions, type `date()` .. The result from what you entered should make it look like

```
Erlang (JAM) emulator version 4.6
```

```
Eshell V4.6 (abort with ^G)
```

```
1> 1 + 2.
```

```
3
```

```
2> date().
```

```
{1998,1,22}
```

The example above can be regarded as having a dialog with the runtime system. You can also run functions or programs, send messages or monitor processes.

To stop Erlang, enter one of the following commands in the window where Erlang was started:

1. Enter the command `C-c` (press the *Control* key and the letter *c* simultaneously). This command puts Erlang into BREAK mode, and the system responds as follows:

```
BREAK: (a)abort (c)ontinue (p)roc info (i)nfo (l)oaded
(v)ersion (k)ill (D)b-tables (d)istribution
```

Enter *a* and then press the *Return* key to abort Erlang.
2. Enter the command `halt()`. and then press the *Return* key.
3. Enter the command `C-g` and then press the *Return* key. This command returns with `->` press *q* to quit and then press the *Return* key to abort Erlang.

You can start various tools from the toolbar



Figure 1.1: The Erlang Toolbar

Compiling a File

This section describes how to compile and run a simple Erlang program, using the Erlang shell. It is assumed that the following program, which calculates the factorial of an integer, has been created with an external editor and saved with the file name `math1.erl`.

Make sure that you are in the directory, where the program file is stored. Use the command `pwd()` to check this and, if needed, change directory by using the command `cd(Dir) ..`

```
-module(math1).
-export([factorial/1]).
factorial(0) -> 1;
factorial(N) -> N * factorial(N-1).
```

Follow these steps to compile the program:

1. Enter the command `c(math1) ..`

```
1> c(math1).
{ok,math1}
2>
```

2. The system returns `{ok, math1}` if the compilation was successful. If unsuccessful, the compiler prints an error message in the following format:

```
math1.erl: line number: error message
```

Refer to the section [Compiling](#) [page 14] in this chapter for more information about compiling code.

Evaluating an Expression

When a module has been successfully compiled, the exported functions in the module can be evaluated by entering the expression at the Erlang shell prompt, followed by a full stop and a carriage return.

The following example illustrates an evaluation, using the `math1.erl` program:

- Enter the following expression at the command prompt:

```
math1:factorial(45).
```

- Erlang evaluates the expression and prints the result. The result from the expression entered looks as follows:

```
2> math1:factorial(45).
119622220865480194561963161495657715064383733760000000000
3>
```

Note:

All Erlang expressions must end with a full stop.

Aborting Evaluation of an Expression

Enter *C-c*. to abort an evaluation (press the *Control* key and the letter *c* simultaneously). This action will also terminate an infinite loop.

How to get help

The Erlang language itself is described in the book [erlbook2], which is available online in PDF² format.

The complete documentation is part of the release in HTML format. The top of the documentation tree is located in the *doc* directory of the Erlang installation. It is also available on-line³

From an Unix shell you can access the manual pages for the Erlang libraries. As an example to display the manual page for the module titled `lists` you enter

```
% erl -man lists
```

You can list the shell commands from the Erlang shell if you enter

```
1> help().
```

The Erlang Shell

The Erlang Shell is a program, which provides a front-end to the Erlang evaluator. Edit commands, shell internal commands, and Erlang expressions, which are entered in the shell window, are interpreted by the Erlang evaluator.

In the shell it is possible to:

- evaluate expressions
- recall previous expressions
- create or destroy variable bindings
- manipulate concurrent jobs.

Depending on your preferred working environment, start Erlang in one of the following ways:

- On a Unix system you enter the command `erl` at the operating system prompt. This command starts the Erlang runtime system and the Erlang shell.
- On the Windows platform you normally start Erlang from the start menu. You can also enter the command `erl` or `werl` from a DOS box or Command box. Note that starting with `erl` will give you a more primitive Erlang shell than if you start with `werl`, see the *werl* reference manual for details.

²URL: <http://www.erlang.org/download/erlang-book-part1.pdf>

³URL: <http://www.erlang.se/erlang/sure/main/products/doc/doc/>

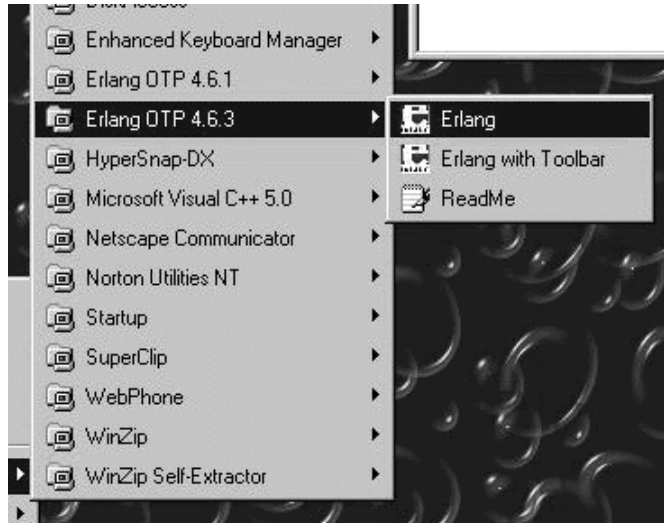


Figure 1.2: Erlang menu on Windows

Apart from evaluation commands and expression, the shell has the following properties:

- Bindings made to variables are remembered.
- Changes made to the local process dictionary (`get` and `put`) are remembered.
- New bindings and changes to the local process dictionary are lost if an expression is entered and its evaluation fails. In this case, the old values are restored.
- Jobs can be manipulated in the shell. A job can be thought of as a set of processes, which can communicate with the shell.

The shell can run in two separate modes:

- *Normal mode*, in which commands can be edited and expressions evaluated
- *Job control mode (JCL)*, in which jobs can be started, stopped, detached or connected. Only the current job can communicate with the shell.

Editing Commands

The shell supports a sub-set of the *Emacs* line editing commands. These commands are listed in the following table.

Note:

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

The *Esc* key is also called the *Meta* key in the Unix environment.

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

Table 1.1: Text Editing

Shell Internal Commands

The shell has a number of built-in internal commands. The command *help()* displays the following commands:

```
1> help().
** shell internal commands **
b()          -- display all variable bindings
e(N)         -- repeat the expression in query <N>
f()          -- forget all variable bindings
h()          -- history
v(N)         -- use the value of query <N>
** commands in module c **
c(File)      -- compile and load code in <File>
cd(Dir)      -- cd Dir .. example cd(".")
flush()      -- flush any messages sent to the shell
help()       -- help info
i()          -- information about the system
ni()         -- information about the networked system
i(X,Y,Z)     -- information about pid <X,Y,Z>
l(File)      -- Load module in <File>
lc([File])  -- Compile a list of Erlang modules
ls()         -- list files in the current directory
ls(Dir)      -- list files in directory <Dir>
m()          -- which modules are loaded
m(Mod)       -- information about module <Mod>
nc(File)     -- compile and load code in <File> on all nodes
nl(File)     -- Load module in <File> on all nodes
pid(X,Y,Z)   -- convert X,Y,Z to a Pid
```

```

pwd()          -- print working directory
regs()        -- information about registered processes
nregs()       -- information about all registered processes
** commands in module i (interpreter interface) **
ih()          -- print help for the i module
true
2>

```

When an expression with the format `Func(Arg1, Arg2, . . . , ArgN)` is entered, the shell has to determine what `Func` refers to before the expression can be evaluated. It interprets `Func` in the following order of priority:

1. `Func` refers to a functional object (Fun)
2. `Func` refers to a built in function (BIF)
3. `user_default:Func` in the module `user_default`
4. `shell_default:Func` in the module `shell_default`

Example Dialogue with the Shell

The following is an extended dialogue with the shell. Detailed explanations of the individual commands are included.

```

host_prompt> erl
Eshell V4.4 (abort with ^G)
1> Str = "abcd".
"abcd"

```

Command 1 sets the variable `Str` to the string "abcd". Remember that ". . ." is the shorthand for the list of ASCII integers which represents the string.

```

2> L = length(Str).
4

```

Command 2 matches the pattern `L` against the return value of `length(Str)`. Since `L` is unbound the match succeeds and `L` is bound to 4.

```

3> Descriptor = {L, list_to_atom(Str)}.
{4,abcd}

```

Command 3 binds the value of variable `Descriptor` to the tuple `{4,abcd}`.

```

4> L.
4

```

Command 4 prints the value of the variable `L`.

```
5> b().
Descriptor = {4,abcd}
L = 4
Str = "abcd"
ok
```

Command 5 evaluates the shell internal command `b()` (short for bindings). This prints a list of the current shell variables and their bindings.

```
6> f(L).
ok
```

Command 6 evaluates the shell internal command `f(L)` (short for forget) which forgets the value of the variable `L`.

```
7> b().
Descriptor = {4,abcd}
Str = "abcd"
ok
```

Command 7 returns a list of the current variable bindings. The variable `L` has been forgotten.

```
8> {L, _} = Descriptor.
{4,abcd}
```

Command 8 performs a pattern matching operation on `Descriptor`, binding a new value to `L`.

```
9> L.
4
```

Command 9 evaluates `L` and prints its value.

```
10> {P, Q, R} = Descriptor.
** exited: {{badmatch,{4,abcd},{erl_eval,expr,3}} **
```

Command 10 tries to match `{P, Q, R}` against `Descriptor` which is `{4, abc}`. The match fails and none of the new variables are bound. `** exited: badmatch ... **` is printed. This is not the value of the expression, and the expression has no value since its evaluation failed. It is a warning printed by the system to inform you that an error has occurred. The information printed after the word `badmatch` indicates which term was being matched, and in which function the error occurred. The values of all other variables, such as `L` and `Str` are unchanged.

```
11> P.
** exited: {{unbound, 'P'},{erl_eval, expr, 3}} **
```

```
12> Descriptor.  
{4,abcd}
```

Commands 11 and 12 show that P is unbound, as the previous command failed, and Descriptor has not changed.

```
13> {P, Q} = Descriptor.  
{4,abcd}
```

```
14> P.  
4
```

Commands 13 and 14 show successful matches, in which P and Q are bound.

```
15> f().  
ok.
```

Command 15 clears all bindings.

For the next commands we assume that `test1:demo(X)` is defined in the following way:

```
test1:demo(X) ->  
put(aa, worked),  
X = 1,  
X + 10.
```

```
16> put(aa, hello).  
undefined  
17> get(aa).  
hello
```

Commands 16 and 17 set and inspect the value of the item aa in the process dictionary.

```
18> Y = test1:demo(1).  
11  
19> get().  
[{aa, worked}]
```

Command 18 evaluates `test1:demo(1)`, the evaluation succeeds and the changes made in the process dictionary become visible to the shell. The new value of the dictionary item aa can be seen in command 19.

```
20> put(aa, hello).
worked
21> Z = test1:demo(2).
!!! Error in process <0.17.0> with exit value:
{{badmatch,1},{test1,demo,[2]}}
** exited: {{badmatch,1},{test1,demo,[2]}} **
```

Commands 20 and 21 reset the value of the dictionary item `aa` to `hello` and calls `test1:demo(2)`. The evaluation fails, and the changes made to the dictionary in `test1:demo(2)` before the error occurred are discarded.

```
22> Z.
** exited: {{unbound,'Z'},{erl_eval,expr,3}} **
23> get(aa).
hello
```

Commands 22 and 23 show that `Z` was not bound and that the dictionary item `aa` retained its original value.

```
24> erase(), put(aa, hello).
undefined.
25> spawn(test1, demo, [1]).
<0.15.1>
26> get(aa).
hello
```

Commands 24, 25, and 26 show the effect of evaluating `test1:demo(1)` in the background. In this case the expression is evaluated in a newly spawned evaluator and any changes made in the process dictionary are local to the newly spawned process and therefore not visible to the shell.

```
27> io:format("hello hello\n").
hello hello
ok
28> e(27).
hello hello
ok
29> v(27).
ok
```

Commands 27, 28, and 29 use the history facilities of the shell. Command 28 is `e(27)`, which re-evaluates command 27. Command 29 is `v(27)` which uses the value (result) of command 27. In the case of a 'pure' function, i.e. a function with no side effects, the result is the same. For a function with side effects the result can be different.

In the next command it is assumed that `test1:loop(N)` is defined in the following way:

```
loop(N) ->
io:format('Hello Number: ~w\n', [N]),
loop(N+1).
```

We continue with command 30.

```
30> test1:loop(0).
Hello Number:1
Hello Number:2
Hello Number:3
Hello Number:4
Hello Number:5
User switch command
-->
```

Command 30 evaluates `test1:loop(0)`, which puts the system into an infinite loop. The user typed Control G. This suspends output from the current process, which is in a loop, and enters JCL mode. In JCL mode you can start and stop jobs.

```
31> F = fun(X) -> 2*X end.
#Fun<erl_eval>
32> F(2).
4
```

Command 31 binds the variable `F` to a *fun*, and command 32 evaluates the expression `F(2)`.

Advanced Shell Usage - Job Control Mode

When the shell starts, it starts a single evaluator process. This process, together with any local processes, which it spawns, is referred to as a *job*. Only the current job, which is said to be *connected*, can perform operations using standard I/O. All other jobs, which are said to be *detached*, will be *blocked* if they attempt to use standard I/O. Jobs which do not use standard I/O run in the normal way.

Typing Control G detaches the current job, and the shell enters into JCL mode. The prompt changes and looks as follows:

```
-->
```

Typing `?` at the prompt gives the following help message:

```
--> ?
c [nn]  - connect to the current job
i [nn]  - interrupt job
k [nn]  - kill job
j       - list all jobs
s       - start new job
r [node] - start a remote shell
q       - quit erlang
? | h   - help
```

The *JCL* commands have the following meaning.

<i>Command</i>	<i>Function</i>
<code>c</code>	Connect to the current job. The standard shell will be resumed. Operations, which use standard I/O, will be interleaved with user inputs to the shell.
<code>c [nn]</code>	As above, but only for job number <code>nn</code> .
<code>k</code>	Kill the current job. All spawned processes in the job will be killed provided they have not evaluated the <code>group_leader/2</code> primitive, and provided they are located on the local machine. Processes spawned on remote nodes will not be killed.
<code>k [nn]</code>	As above but specifically only for job number <code>nn</code> .
<code>j</code>	List all jobs. A list of all known jobs is printed. The current job name is prefixed with <code>**</code> .
<code>s</code>	Start a new job. This will be assigned a new index <code>[nn]</code> , which can be used for referencing.
<code>r [node]</code>	Start a remote job on <code>node</code> . This is used in distributed Erlang to allow a shell running on one node to control a number of applications running on a network of nodes.
<code>q</code>	Quit Erlang.
<code>? h</code>	Help.

Table 1.2: The JCL Commands

Compiling

A file is compiled by evaluating the function

```
compile:file(File, Options).
```

`File` is a file name, and `Options` is a list of compiler options. Refer to the Reference Manual for a complete list of `Options`.

You can also enter the command `c(File)` from the Erlang shell. This command assumes that the following conditions apply:

1. The module name is the same as the file name, minus the `.erl` extension.
2. Code is to be loaded immediately if the compilation succeeds.

The function `make:all()` can be used to compile many files. All files must be located in the same directory. Refer to the Reference Manual, function `make(3)` for full details.

Compiler Diagnostics

The Erlang compiler detects a number of different errors. These can be:

- syntax errors
- variable Errors
- missing function errors.

The compiler also warns about potential errors.

Syntax Errors Syntax errors are reported against syntactically incorrect Erlang code. For example, the following code fragment is incorrect and will be reported by the compiler as shown:

```
test(X, Y) ->
  case g(X) o
    Y -> 1;
    2 -> 2
  end.
```

The syntax word `of` in the case construction is incorrectly spelt. This produces the compiler diagnostic:

```
test1.erl: 10 : of expected before 'o'
```

This means that an error occurred in line 10 of the file `test1.erl`. The message which, follows the line number, describes the error.

Variable Errors Two types of variable errors can be detected by the compiler. Variables can be *undefined*, or *unsafe*. In the following example, the variables `Y` and `Z` are undefined:

```
test(X) ->
  A = f(X, Y, Z),
  b(A, X).
```

When compiled, the following error messages are printed:

```
test.erl: 11 : unbound variable 'Y'
test.erl: 11 : unbound variable 'Z'
```

The reason for these errors are that the variables `Y` and `Z` are undefined.

The following example shows the use of unsafe variables:

```
test1(X) ->
  case f(X) of
    true -> X = 1, Y=2;
    false -> X = 2
  end,
  g(X, Y).
```

In this example, there are two possible execution paths in the case statement. In one case the variables `X` and `Y` are bound, in the other case only the variable `X` is bound. After the case statement, we can only be certain that `X` is bound and the variable `Y` is said to be *unsafe*. Any subsequent use of `Y` will result in an unsafe variable error message. In this case the following compiler error is generated:

```
test1.erl : 15 unsafe variable 'Y'
```

If no reference to `Y` is made after the case statement, the function is correct, as the following example shows:

```
test2(X) ->
  case f(X) of
    true -> X = 1, Y=2;
    false -> X = 2
  end,
  g(X).
```

Missing Functions The following module contains a missing function:

```
-module(test).
-export([test/1]).
test(X) -> try(X, abc).
```

When the module is compiled, it results in the following:

```
> c(test).
test.erl : 25 : Missing function test:try/2
```

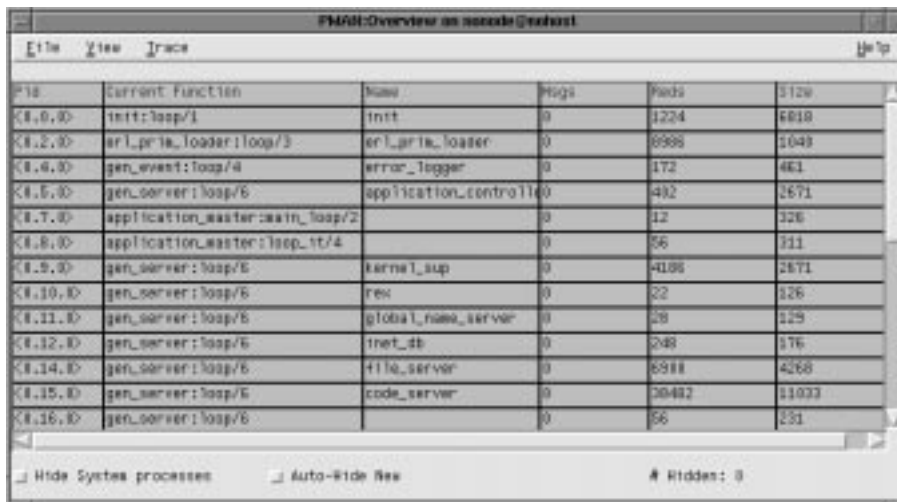
Debugging with the Process Manager

The Process Manager, *pman*, is a tool which traces the evaluation of code. It can trace several different processes, and several different machines if distributed Erlang is used. *pman* is controlled through a simple graphical user interface.

Start *pman* as follows:

- run the function `pman:start()` from the Erlang shell.
- push the debugger button on the Toolbar.

The following window is displayed when starting the *pman*:



The screenshot shows a window titled "PMan-Overview on oscode@oscode". It contains a table with the following columns: PID, Current Function, Name, Pugs, Pids, and Size. The table lists several processes, including 'init', 'erl_prt_loader', 'error_logger', 'application_control', 'application_master', 'application_master', 'kernel_sup', 'rex', 'global_name_server', 'inet_db', 'file_server', 'code_server', and another 'application_master'.

PID	Current Function	Name	Pugs	Pids	Size
<1.0.0>	init:loop/1	init	0	1224	6818
<1.2.0>	erl_prt_loader:loop/3	erl_prt_loader	0	8988	1043
<1.4.0>	gen_event:loop/4	error_logger	0	172	461
<1.5.0>	gen_server:loop/6	application_control	0	432	2671
<1.7.0>	application_master:main_loop/2	application_master	0	12	326
<1.8.0>	application_master:loop_it/4	application_master	0	56	311
<1.9.0>	gen_server:loop/6	kernel_sup	0	4086	2671
<1.10.0>	gen_server:loop/6	rex	0	22	126
<1.11.0>	gen_server:loop/6	global_name_server	0	28	129
<1.12.0>	gen_server:loop/6	inet_db	0	248	176
<1.14.0>	gen_server:loop/6	file_server	0	6918	4268
<1.15.0>	gen_server:loop/6	code_server	0	30432	11033
<1.16.0>	gen_server:loop/6	application_master	0	56	231

Figure 1.3: The pman Window

This section describes a simple debugging example, which involves tracing the code in a single process. For a more detailed description of *pman* see the User's Manual.

Tracing a Single Function

We illustrate the use by tracing the evaluation of the function `lists1:reverse`, which is defined as follows:

```
-module(lists1).  
-export([reverse/1]).  
reverse(X) -> reverse(X, []).  
reverse([H|T], L) -> reverse(T, [H|L]);  
reverse([], L) -> L.
```

Follow these steps to trace this process:

1. Select the *Shell Process* item in the *Trace* menu of the pman main window. The window shown in [The Shell Tracer Window \[page 18\]](#) is displayed.
2. From the shell tracer window select the *Options...* item in the *File* menu. Select the *Trace functions* check box and click *OK*.
3. The first command `c(lists1, trace)` compiles the module `lists1` in trace mode.

```
1> c(lists1, trace).  
{ok, lists1}
```

4. Evaluate the function `lists1:reverse("abcd")`

```
2> lists1:reverse("abcd").  
"dcba".
```

5. The contents of the shell trace window changes to show a listing of the called functions. See [Trace of the function lists1:reverse \[page 18\]](#).

Note:

All traced modules must be compiled in trace mode.

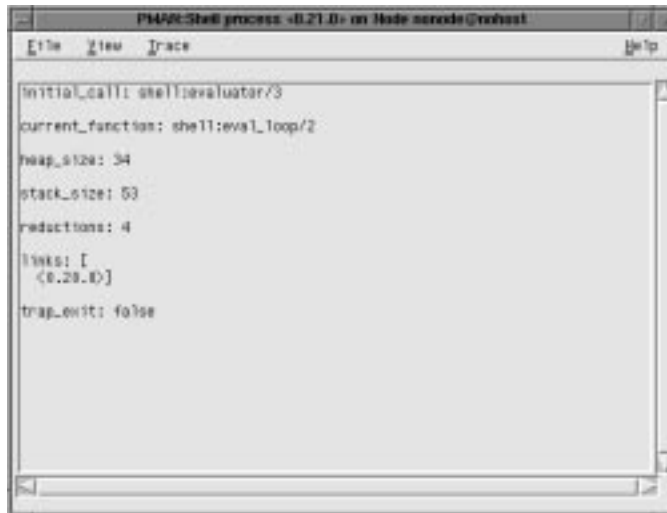


Figure 1.4: The Shell Tracer Window




Figure 1.5: Trace of the Function `lists:reverse`

Distributed Erlang

Erlang can be used to write distributed applications, which run simultaneously on several computers. Programming distributed systems with Erlang is the subject of the book [erlbook2]. Refer also to the Reference Manual.

Nodes

A *node* is a basic concept of distributed Erlang. In a distributed Erlang system, the term node means an executing Erlang system, which can communicate with other Erlang systems. Each node has a name assigned to it, which is returned by the built in function `node()`. This name is an atom and is guaranteed to be globally unique. The format of the name can vary between different implementations of Erlang, but it is always an atom which consists of two parts separated by an `@` character. The following is an example of a node name:

```
dilbert@super.du.ericsson.se
```

Before we can fully account for the structure of a node name, we have to consider how host names are specified.

Host Names Every host, which is attached to a network, has an IP address. However, it is not practical to use IP addresses when referring to hosts, since they are made up of numeric strings, which are difficult to remember. For this reason, hosts are usually given names.

Host names of hosts, which are connected to the Internet, are organized into *domains*. Domains are built hierarchically, with “top level” domains and “sub-domains”. For example, the domain denoted by `du.etx.ericsson.se` is a sub-domain, which is part of another sub-domain `etx.ericsson.se`, which is part of the top domain `se`.

When it is clear from the context what domain a host belongs to, it is not necessary to use the fully qualified host name. In the example `super.du.etx.ericsson.se`, the first part “super” will suffice. We refer to “super” as the *short host name* of the host.

Hosts which are not connected to the Internet, and do not belong to any particular domain, are simply referred to by their short host name, “super” for example.

We use the term *host name* to denote the name of hosts. It should be clear from the context if we mean the long or the short form of the name.

Node Names An Erlang node name is an atom constructed as the concatenation of a name supplied by the user, the `@` character, and the name of the host where the Erlang system is executed. For example:

```
dilbert@super.du.etx.ericsson.se
```

In this example, “dilbert” is the name of the Erlang node, and “super.du.etx.ericsson.se” is the name of the host where Erlang is running.

Distributed Erlang requires host names to be unique, since Erlang nodes identify each other by node names. We can achieve uniqueness for host names either by using long host names, or using short host names with the additional requirement that communicating hosts belong to the same domain (or a network not connected to Internet).

When using long host names, the corresponding node names are called *long node names*. For example:

```
dilbert@super.du.etx.ericsson.se
```

In this case, distributed Erlang is started with the `-name` flag. The system checks that the host name part of the node name contains at least one dot (.). If not, the host name is considered to be in error.

Similarly, when using short host names, the corresponding node names are called *short node names*. For example:

```
dilbert@super
```

In this case, distributed Erlang is started with the `-sname` flag.

Note:

An Erlang started with a long node name cannot communicate with a node started with a short node name.

It is often clear from the context if a node name is long or short. When this is the case, we use the term *node name* to refer to a node (the name of which may be short or long).

The first part of a node name which comes before the `@` character is the *plain name* of the node. The plain name for a node is unique within the host on which the node is running. When starting distributed Erlang, with either the `-name` or the `-sname` flag, the plain name of the node is provided by the user. The other part of the node name, which is the host name part, is added by the Erlang system at start-up.

Start-Up

Execute the following command to start a distributed Erlang node:

```
% erl -name foobar
Eshell V4.4 (abort with ^G)
(foobar@super.eua.ericsson.se)1>
```

The node name shown in the system response, includes the long name for the host. An alternative way is to start the system with the `-sname` flag. This is sometimes the only way to run distributed erlang if the Domain Name System (DNS) is not running, or is not available.

Distribution Script Flags

The executable script `erl` starts the Erlang runtime system. Refer to the Reference Manual, `erl(1)`, for a complete list of the flags, which can be added to the start script. This section describes some of the flags, which control the distribution:

- `-name Name` This flag makes the node into a distributed node. It invokes all network servers which are needed to make a node a distributed node. Refer to the Reference Manual, `net_kernel(3)` for full details. The name of the node will be `Name@Host`, where `Host` is the fully qualified host name of the current host. This option will also ensure that the Erlang port mapper daemon (`epmd`) runs on the current host before Erlang is started. Refer to the Reference Manual, `epmd(3)`.

- `-sname Name` This flag is the same as the `-name` flag, except for the host name used. The name of the node will be `Name@Host`, where `Host` is the short host name of the current host. If the Domain Name System (DNS) is not running, this is sometimes the only way to run distributed Erlang Systems. The `-sname` flag cannot communicate with systems running with the `-name` flag, since node names must be unique in distributed Erlang.
- `-setcookie Cookie` This flag sets the magic cookie of the current node to `Cookie`. Since `erlang:set_cookie(node(), Cookie)` is used, all other nodes are assumed to have their cookies set to `Cookie` as well. This way several nodes can share a common magic cookie (see the next section Security [page 21]).
- `-connect_all false` This flag causes the node connection graph to become 'leaner'. Only direct connections are established.

Security

Erlang nodes are protected by a magic cookie system. Refer to the Reference Manual, section `kernel`, the module `auth` (3) for detailed information about the magic cookie system.

1. When a message is sent from one node to another node, the sending node attaches the magic cookie of the receiving node to the message.
2. The runtime system on the receiving node verifies that the magic cookie is the correct one for this node.
3. If the magic cookie is correct, the message is accepted. If the magic cookie is not correct, then the message is handled as follows:
 - The message is transformed into a "badcookie" message.
 - The "badcookie" message is sent to the system process `net_kernel`.
 - By default, the `net_kernel` passes the message to the registered process `auth`.
 - The `auth` process is then responsible for taking appropriate action for the unauthorized message. In the standard system, the default action is to shut down the connection to the sending node.

Magic cookies are assigned to a node as follows:

1. Initially, each node is assigned a random atom as its magic cookie, and each node assumes the atom `nocookie` for all other nodes.
2. When the standard `auth` server is started without the `-setcookie` flag, it reads a file named `$HOME/.erlang.cookie`. If the file `$HOME/.erlang.cookie` does not exist, it is created with a random string as its content. The permission mode of the file is set to `read-only` by owner.
3. The system creates an atom from the contents of this file and the function `erlang:set_cookie(node(), CookieAtom)` sets the magic cookie for the node to the value of this atom.

An Erlang system is completely unprotected when running `erlang:set_cookie(node(), nocookie)`. This can sometimes be appropriate for systems that are not normally networked, or for systems which are run for maintenance purposes only. Refer to the Reference Manual, `auth(3)` for details on the security system.

Slave Nodes

The Erlang library module `slave` provides functionality to start or control slave nodes. This functionality is Unix specific.

The remote nodes are started by means of the Unix command `rsh`. However, to use this command, the user must not only use a Unix system, but must also be allowed to `rsh` to the remote hosts without being prompted for a password.

Refer to the following manuals for additional information:

- The Unix documentation set, `rsh(1)`
- The Reference Manual, `slave(3)`.

Configuration

The standard Erlang system can be re-configured to change the default behavior on start-up.

The `.erlang` Start-up File

When Erlang is started, the system searches for a file named `.erlang` in the directory where Erlang is started. If not found, the user's home directory is searched for an `.erlang` file.

If an `.erlang` file is found, it is assumed to contain valid Erlang expressions. These expressions are evaluated as if they were input to the shell.

A typical `.erlang` file contains a set of search paths, for example:

```
io:format("executing user profile in HOME/.erlang\n", []).
code:add_uc(spy).
code:add_uc(pxw).
code:add_uc(string).
code:add_uc(unix).
code:add_uc(pretty_format).
code:add_uc(iv20).
code:add_uc(string).
io:format(".erlang rc finished\n", []).
```

`user_default` and `shell_default`

Functions in the shell which are not prefixed by a module name are assumed to be functional objects (Funs), built in functions (BIFs), or belong to the module `user_default` or `shell_default`.

To include private commands, define them in a module `user_default` and add the following argument as the first line in the `.erlang` file.

```
code:load_abs("../user_default)
```

`erl`

If the contents of `.erlang` are changed and a private version of `user_default` is defined, it is possible to customize the Erlang environment. More powerful changes can be made by supplying command line arguments in the start-up script `erl`. Refer to the Reference Manual, `erl(1)` and `init(3)` for further information.

List of Figures

Chapter 1: Getting Started

1.1	The Erlang Toolbar	4
1.2	Erlang menu on Windows	7
1.3	The pman Window	16
1.4	The Shell Tracer Window	18
1.5	Trace of the Function lists1:reverse	18

List of Tables

Chapter 1: Getting Started

- 1.1 Text Editing 8
- 1.2 The JCL Commands 14