

SBCL User Manual

SBCL version 2.1.11-bp154.1.169-suse
2022-05

This manual is part of the SBCL software system. See the `README` file for more information.

This manual is largely derived from the manual for the CMUCL system, which was produced at Carnegie Mellon University and later released into the public domain. This manual is in the public domain and is provided with absolutely no warranty. See the `COPYING` and `CREDITS` files for more information.

Table of Contents

1	Getting Support and Reporting Bugs	1
1.1	Volunteer Support	1
1.2	Commercial Support	1
1.3	Reporting Bugs	1
1.3.1	How to Report Bugs Effectively	1
1.3.2	Signal Related Bugs	2
2	Introduction	3
2.1	ANSI Conformance	3
2.1.1	Exceptions	3
2.2	Extensions	3
2.3	Idiosyncrasies	4
2.3.1	Declarations	4
2.3.2	FASL Format	4
2.3.3	Compiler-only Implementation	5
2.3.4	Defining Constants	5
2.3.5	Style Warnings	5
2.4	Development Tools	5
2.4.1	Editor Integration	6
2.4.2	Language Reference	6
2.4.3	Generating Executables	6
2.5	More SBCL Information	6
2.5.1	SBCL Homepage	6
2.5.2	Online Documentation	6
2.5.3	Additional Documentation Files	6
2.5.4	Internals Documentation	6
2.6	More Common Lisp Information	7
2.6.1	Internet Community	7
2.6.2	Third-party Libraries	7
2.6.3	Common Lisp Books	7
2.7	History and Implementation of SBCL	7
3	Starting and Stopping	9
3.1	Starting SBCL	9
3.1.1	From Shell to Lisp	9
3.1.2	Running from Emacs	9
3.1.3	Shebang Scripts	9
3.2	Stopping SBCL	9
3.2.1	Exit	9
3.2.2	End of File	10
3.2.3	Saving a Core Image	10
3.2.4	Exit on Errors	12
3.3	Command Line Options	12
3.3.1	Runtime Options	12
3.3.2	Toplevel Options	13
3.4	Initialization Files	14
3.5	Initialization and Exit Hooks	14
4	Compiler	15
4.1	Diagnostic Messages	15
4.1.1	Controlling Verbosity	15
4.1.2	Diagnostic Severity	16

4.1.3	Understanding Compile Diagnostics	16
4.1.3.1	The Parts of a Compiler Diagnostic	16
4.1.3.2	The Original and Actual Source	18
4.1.3.3	The Processing Path	18
4.2	Handling of Types	19
4.2.1	Declarations as Assertions	19
4.2.2	Precise Type Checking	19
4.2.3	Getting Existing Programs to Run	20
4.2.4	Implementation Limitations	21
4.3	Compiler Policy	21
4.4	Compiler Errors	23
4.4.1	Type Errors at Compile Time	23
4.4.2	Errors During Macroexpansion	24
4.4.3	Read Errors	24
4.5	Open Coding and Inline Expansion	24
4.6	Interpreter	25
4.7	Advanced Compiler Use and Efficiency Hints	25
5	Debugger	26
5.1	Debugger Entry	26
5.1.1	Debugger Banner	26
5.1.2	Debugger Invocation	26
5.2	Debugger Command Loop	26
5.3	Stack Frames	27
5.3.1	Stack Motion	27
5.3.2	How Arguments are Printed	27
5.3.3	Function Names	28
5.3.3.1	Entry Point Details	28
5.3.4	Debug Tail Recursion	28
5.3.5	Unknown Locations and Interrupts	29
5.4	Variable Access	29
5.4.1	Variable Value Availability	30
5.4.2	Note On Lexical Variable Access	30
5.5	Source Location Printing	31
5.5.1	How the Source is Found	31
5.5.2	Source Location Availability	32
5.6	Debugger Policy Control	32
5.7	Exiting Commands	33
5.8	Information Commands	33
5.9	Function Tracing	34
5.10	Single Stepping	35
5.11	Enabling and Disabling the Debugger	36
6	Efficiency	37
6.1	Slot access	37
6.1.1	Structure object slot access	37
6.1.2	Standard object slot access	37
6.2	Dynamic-extent allocation	37
6.3	Modular arithmetic	39
6.4	Global and Always-Bound variables	39
6.5	Miscellaneous Efficiency Issues	40

7	Beyond the ANSI Standard	42
7.1	Reader Extensions	42
7.1.1	Extended Package Prefix Syntax	42
7.1.2	Symbol Name Normalization	42
7.1.3	Decimal Syntax for Rationals	42
7.2	Package-Local Nicknames	42
7.3	Package Variance	43
7.4	Garbage Collection	44
7.4.1	Finalization	44
7.4.2	Weak Pointers	45
7.4.3	Introspection and Tuning	45
7.4.4	Tracing Live Objects Back to Roots	46
7.5	Metaobject Protocol	47
7.5.1	AMOP Compatibility of Metaobject Protocol	47
7.5.2	Metaobject Protocol Extensions	49
7.6	Extensible Sequences	49
7.6.1	Iterator Protocol	52
7.6.2	Simple Iterator Protocol	52
7.7	Support For Unix	53
7.7.1	Command-line arguments	53
7.7.2	Querying the process environment	53
7.7.3	Running external programs	53
7.8	Unicode Support	56
7.8.1	Unicode property access	56
7.8.2	String operations	58
7.8.3	Breaking strings	59
7.9	Customization Hooks for Users	59
7.10	Tools To Help Developers	60
7.11	Resolution of Name Conflicts	60
7.12	Hash Table Extensions	60
7.13	Random Number Generation	62
7.14	Timeouts and Deadlines	63
7.14.1	Timeout Parameters	63
7.14.2	Synchronous Timeouts (Deadlines)	64
7.14.3	Asynchronous Timeouts	64
7.14.4	Operations Supporting Timeouts and Deadlines	65
7.15	Miscellaneous Extensions	65
7.16	Stale Extensions	66
7.17	Efficiency Hacks	66
8	External Formats	67
8.1	The Default External Format	67
8.2	External Format Designators	67
8.3	Character Coding Conditions	67
8.4	Converting between Strings and Octet Vectors	68
8.5	Supported External Formats	68
9	Foreign Function Interface	71
9.1	Introduction to the Foreign Function Interface	71
9.2	Foreign Types	71
9.2.1	Defining Foreign Types	71
9.2.2	Foreign Types and Lisp Types	72
9.2.3	Foreign Type Specifiers	72
9.3	Operations On Foreign Values	74
9.3.1	Accessing Foreign Values	74
9.3.1.1	Untyped memory	74
9.3.2	Coercing Foreign Values	74

9.3.3	Foreign Dynamic Allocation	75
9.4	Foreign Variables	75
9.4.1	Local Foreign Variables	75
9.4.2	External Foreign Variables	76
9.5	Foreign Data Structure Examples	76
9.6	Loading Shared Object Files	77
9.7	Foreign Function Calls	78
9.7.1	The <code>alien-funcall</code> Primitive	78
9.7.2	The <code>define-alien-routine</code> Macro	78
9.7.3	<code>define-alien-routine</code> Example	79
9.8	Calling Lisp From C	79
9.8.1	Lisp as a Shared Library	80
9.9	Step-By-Step Example of the Foreign Function Interface	80
10	Pathnames	83
10.1	Lisp Pathnames	83
10.1.1	Home Directory Specifiers	83
10.1.2	The SYS Logical Pathname Host	83
10.2	Native Filenames	83
11	Streams	85
11.1	Stream External Formats	85
11.2	Bivalent Streams	85
11.3	Gray Streams	85
11.3.1	Gray Streams classes	85
11.3.2	Methods common to all streams	86
11.3.3	Input stream methods	86
11.3.4	Character input stream methods	87
11.3.5	Output stream methods	87
11.3.6	Character output stream methods	87
11.3.7	Binary stream methods	88
11.3.8	Gray Streams examples	88
11.3.8.1	Character counting input stream	88
11.3.8.2	Output prefixing character stream	90
11.4	Simple Streams	91
12	Package Locks	92
12.1	Package Lock Concepts	92
12.1.1	Package Locking Overview	92
12.1.2	Implementation Packages	92
12.1.3	Package Lock Violations	92
12.1.3.1	Lexical Bindings and Declarations	92
12.1.3.2	Other Operations	92
12.1.4	Package Locks in Compiled Code	93
12.1.4.1	Interned Symbols	93
12.1.4.2	Other Limitations on Compiled Code	93
12.1.5	Operations Violating Package Locks	93
12.1.5.1	Operations on Packages	93
12.1.5.2	Operations on Symbols	93
12.2	Package Lock Dictionary	94

13	Threading	96
13.1	Threading basics	96
13.1.1	Thread Objects	96
13.1.2	Making, Returning From, Joining, and Yielding Threads	96
13.1.3	Asynchronous Operations	97
13.1.4	Miscellaneous Operations	98
13.1.5	Error Conditions	98
13.2	Special Variables	99
13.3	Atomic Operations	99
	CAS Protocol	101
13.4	Mutex Support	101
13.5	Semaphores	103
13.6	Waitqueue/condition variables	103
13.7	Barriers	105
13.8	Sessions/Debugging	106
13.9	Foreign threads	106
13.10	Implementation (Linux x86/x86-64)	106
14	Timers	108
14.1	Timer Dictionary	108
15	Networking	110
15.1	Sockets Overview	110
15.2	General Sockets	110
15.3	Socket Options	112
15.4	INET Domain Sockets	112
15.5	Local (Unix) Domain Sockets	113
15.6	Name Service	113
16	Profiling	114
16.1	Deterministic Profiler	114
16.2	Statistical Profiler	114
16.2.1	Example Usage	114
16.2.2	Output	115
16.2.3	Platform support	116
16.2.4	Macros	116
16.2.5	Functions	117
16.2.6	Variables	118
16.2.7	Credits	118
17	Contributed Modules	119
17.1	sb-aclrepl	120
17.1.1	Usage	120
17.1.2	Customization	120
17.1.3	Example Initialization	120
17.1.4	Credits	120
17.2	sb-concurrency	121
17.2.1	Queue	122
17.2.2	Mailbox (lock-free)	123
17.2.3	Gates	124
17.2.4	Frlocks, aka Fast Read Locks	125
17.3	sb-cover	127
17.3.1	Example Usage	127
17.3.2	Functions	127
17.4	sb-grovel	129
17.4.1	Using sb-grovel in your own ASDF system	129

17.4.2	Contents of a grovel-constants-file	129
17.4.3	Programming with sb-grovel's structure types.....	130
17.4.3.1	Traps and Pitfalls	131
17.5	sb-md5.....	132
17.5.1	Credits	132
17.6	sb-posix.....	133
17.6.1	Lisp names for C names	133
17.6.2	Types.....	133
17.6.2.1	File-descriptors	133
17.6.2.2	Filenames.....	133
17.6.3	Function Parameters	134
17.6.4	Function Return Values	134
17.6.5	Lisp objects and C structures	134
17.6.6	Functions with idiosyncratic bindings.....	137
17.7	sb-queue	138
17.8	sb-rotate-byte.....	139
18	Deprecation	140
18.1	Why Deprecate?	140
18.2	The Deprecation Pipeline.....	140
18.3	Deprecation Conditions.....	141
18.4	Introspecting Deprecation Information.....	142
18.5	Deprecation Declaration	142
18.6	Deprecation Examples.....	142
18.7	Deprecated Interfaces in SBCL	143
18.7.1	List of Deprecated Interfaces	143
18.7.1.1	Early Deprecation.....	143
18.7.1.2	Late Deprecation.....	144
18.7.1.3	Final Deprecation	145
18.7.2	Historical Interfaces	145
Appendix A	Concept Index	146
Appendix B	Function Index.....	148
Appendix C	Variable Index	153
Appendix D	Type Index.....	154
Colophon	156

1 Getting Support and Reporting Bugs

1.1 Volunteer Support

Your primary source of SBCL support should probably be the mailing list **sbcl-help**: in addition to other users SBCL developers monitor this list and are available for advice. As an anti-spam measure subscription is required for posting:

<https://lists.sourceforge.net/lists/listinfo/sbcl-help>

Remember that the people answering your question are volunteers, so you stand a much better chance of getting a good answer if you ask a good question.

Before sending mail, check the list archives at either

http://sourceforge.net/mailarchive/forum.php?forum_name=sbcl-help

or

<http://news.gmane.org/gmane.lisp.steel-bank.general>

to see if your question has been answered already. Checking the bug database is also worth it See Section 1.3 [Reporting Bugs], page 1, to see if the issue is already known.

For general advice on asking good questions, see

<http://www.catb.org/~esr/faqs/smart-questions.html>.

1.2 Commercial Support

There is no formal organization developing SBCL, but if you need a paid support arrangement or custom SBCL development, we maintain the list of companies and consultants below. Use it to identify service providers with appropriate skills and interests, and contact them directly.

The SBCL project cannot verify the accuracy of the information or the competence of the people listed, and they have provided their own blurbs below: you must make your own judgement of suitability from the available information - refer to the links they provide, the CREDITS file, mailing list archives, CVS commit messages, and so on. Please feel free to ask for advice on the sbcl-help list.

(At present, no companies or consultants wish to advertise paid support or custom SBCL development in this manual).

1.3 Reporting Bugs

SBCL uses Launchpad to track bugs. The bug database is available at

<https://bugs.launchpad.net/sbcl>

Reporting bugs there requires registering at Launchpad. However, bugs can also be reported on the mailing list **sbcl-bugs**, which is moderated but does *not* require subscribing.

Simply send email to sbcl-bugs@lists.sourceforge.net and the bug will be checked and added to Launchpad by SBCL maintainers.

1.3.1 How to Report Bugs Effectively

Please include enough information in a bug report that someone reading it can reproduce the problem, i.e. don't write

Subject: apparent bug in PRINT-OBJECT (or *PRINT-LENGTH*?)
PRINT-OBJECT doesn't seem to work with *PRINT-LENGTH*. Is this a bug?

but instead

Subject: apparent bug in PRINT-OBJECT (or *PRINT-LENGTH*?)
In sbcl-1.2.3 running under OpenBSD 4.5 on my Alpha box, when
I compile and load the file
(DEFSTRUCT (FOO (:PRINT-OBJECT (LAMBDA (X Y)
 (LET ((*PRINT-LENGTH* 4))
 (PRINT X Y))))))

X Y)

then at the command line type

(MAKE-FOO)

the program loops endlessly instead of printing the object.
A more in-depth discussion on reporting bugs effectively can be found at
<http://www.chiark.greenend.org.uk/~sgtatham/bugs.html>.

1.3.2 Signal Related Bugs

If you run into a signal related bug, you are getting fatal errors such as `signal N is [un]blocked` or just hangs, and you want to send a useful bug report then:

1. Compile SBCL with `ldb` enabled (feature `:sb-ldb`, see `base-target-features.lisp-expr`) and change `#define QSHOW_SIGNAL 0` to `#define QSHOW_SIGNAL 1` in `src/runtime/runtime.h`.
2. Isolate a smallish test case, run it.
3. If it just hangs kill it with `sigabrt`: `kill -ABRT <pidof sbcl>`.
4. Print the backtrace from `ldb` by typing `ba`.
5. Attach `gdb`: `gdb -p <pidof sbcl>` and get backtraces for all threads: `thread apply all ba`.
6. If multiple threads are in play then still in `gdb`, try to get Lisp backtrace for all threads: `thread apply all call backtrace_from_fp($ebp, 100)`. Substitute `$ebp` with `$rbp` on x86-64. The backtraces will appear in the stdout of the SBCL process.
7. Send a report with the backtraces and the output (both stdout and stderr) produced by SBCL.
8. Don't forget to include OS and SBCL version.
9. If available, include information on outcome of the same test with other versions of SBCL, OS, ...

2 Introduction

SBCL is a mostly-conforming implementation of the ANSI Common Lisp standard. This manual focuses on behavior which is specific to SBCL, not on behavior which is common to all implementations of ANSI Common Lisp.

2.1 ANSI Conformance

Essentially every type of non-conformance is considered a bug. (The exceptions involve internal inconsistencies in the standard.) See Section 1.3 [Reporting Bugs], page 1.

2.1.1 Exceptions

- `prog2` returns the primary value of its second form, as specified in the **Arguments and Values** section of the specification for that operator, not that of its first form, as specified in the **Description**.
- The `string` type is considered to be the union of all types (`array c (size)`) for all non-`nil` subtypes `c` of `character`, excluding arrays specialized to the empty type.

2.2 Extensions

SBCL comes with numerous extensions, some in core and some in modules loadable with `require`. Unfortunately, not all of these extensions have proper documentation yet.

System Definition Tool

`asdf` is a flexible and popular protocol-oriented system definition tool by Daniel Barlow. See Info file `asdf`, node ‘Top’, for more information.

Foreign Function Interface

`sb-alien` package allows interfacing with C-code, loading shared object files, etc. See Chapter 9 [Foreign Function Interface], page 71.

`sb-grovel` can be used to partially automate generation of foreign function interface definitions. See Section 17.4 [sb-grovel], page 129.

Recursive Event Loop

SBCL provides a recursive event loop (`serve-event`) for doing non-blocking IO on multiple streams without using threads.

Timeouts and Deadlines

SBCL allows restricting the execution time of individual operations or parts of a computation using `:timeout` arguments to certain blocking operations, synchronous timeouts and asynchronous timeouts. The latter two affect operations without explicit timeout support (such as standard functions and macros). See Section 7.14 [Timeouts and Deadlines], page 63.

Metaobject Protocol

`sb-mop` package provides a metaobject protocol for the Common Lisp Object System as described in *Art of Metaobject Protocol*.

Extensible Sequences

SBCL allows users to define subclasses of the `sequence` class. See Section 7.6 [Extensible Sequences], page 49.

Native Threads

SBCL has native threads on x86/Linux, capable of taking advantage of SMP on multiprocessor machines. See Chapter 13 [Threading], page 96.

Network Interface

`sb-bsd-sockets` is a low-level networking interface, providing both TCP and UDP sockets. See Chapter 15 [Networking], page 110.

Introspective Facilities

`sb-introspect` module offers numerous introspective extensions, including access to function lambda-lists and a cross referencing facility.

Operating System Interface

sb-ext contains a number of functions for running external processes, accessing environment variables, etc.

sb-posix module provides a lispy interface to standard POSIX facilities.

Extensible Streams

sb-gray is an implementation of *Gray Streams*. See Section 11.3 [Gray Streams], page 85.

sb-simple-streams is an implementation of the *simple streams* API proposed by Franz Inc. See Section 11.4 [Simple Streams], page 91.

Profiling

sb-profile is a exact per-function profiler. See Section 16.1 [Deterministic Profiler], page 114.

sb-sprof is a statistical profiler, capable of call-graph generation and instruction level profiling, which also supports allocation profiling. See Section 16.2 [Statistical Profiler], page 114.

Customization Hooks

SBCL contains a number of extra-standard customization hooks that can be used to tweak the behaviour of the system. See Section 7.9 [Customization Hooks for Users], page 59.

sb-aclrepl provides an Allegro CL -style toplevel for SBCL, as an alternative to the classic CMUCL-style one. See Section 17.1 [sb-aclrepl], page 120.

CLTL2 Compatibility Layer

sb-cltl2 module provides **compiler-let** and environment access functionality described in *Common Lisp The Language, 2nd Edition* which were removed from the language during the ANSI standardization process.

Executable Delivery

The **:executable** argument to [Function **sb-ext save-lisp-and-die**], page 10, can produce a ‘standalone’ executable containing both an image of the current Lisp session and an SBCL runtime.

Bitwise Rotation

sb-rotate-byte provides an efficient primitive for bitwise rotation of integers, an operation required by e.g. numerous cryptographic algorithms, but not available as a primitive in ANSI Common Lisp. See Section 17.8 [sb-rotate-byte], page 139.

Test Harness

sb-rt module is a simple yet attractive regression and unit-test framework.

MD5 Sums

sb-md5 is an implementation of the MD5 message digest algorithm for Common Lisp, using the modular arithmetic optimizations provided by SBCL. See Section 17.5 [sb-md5], page 132.

2.3 Idiosyncrasies

The information in this section describes some of the ways that SBCL deals with choices that the ANSI standard leaves to the implementation.

2.3.1 Declarations

Declarations are generally treated as assertions. This general principle, and its implications, and the bugs which still keep the compiler from quite satisfying this principle, are discussed in Section 4.2.1 [Declarations as Assertions], page 19.

2.3.2 FASL Format

SBCL fasl-format is binary compatible only with the exact SBCL version it was generated with. While this is obviously suboptimal, it has proven more robust than trying to maintain fasl compatibility across versions: accidentally breaking things is far too easy, and can lead to hard to diagnose bugs.

The following snippet handles fasl recompilation automatically for ASDF-based systems, and makes a good candidate for inclusion in the user or system initialization file (see Section 3.4 [Initialization Files], page 14.)

```
(require :asdf)
```

```
;;; If a fasl was stale, try to recompile and load (once).
(defmethod asdf:perform :around ((o asdf:load-op)
                                   (c asdf:cl-source-file))
  (handler-case (call-next-method o c)
    ;; If a fasl was stale, try to recompile and load (once).
    (sb-ext:invalid-fasl ()
      (asdf:perform (make-instance 'asdf:compile-op) c)
      (call-next-method))))
```

2.3.3 Compiler-only Implementation

SBCL is essentially a compiler-only implementation of Common Lisp. That is, for all but a few special cases, `eval` creates a lambda expression, calls `compile` on the lambda expression to create a compiled function, and then calls `funcall` on the resulting function object. A more traditional interpreter is also available on default builds; it is usually only called internally. This is explicitly allowed by the ANSI standard, but leads to some oddities; e.g. at default settings, `functionp` and `compiled-function-p` are equivalent, and they collapse into the same function when SBCL is built without the interpreter.

2.3.4 Defining Constants

SBCL is quite strict about ANSI's definition of `defconstant`. ANSI says that doing `defconstant` of the same symbol more than once is undefined unless the new value is `eq` to the old value. Conforming to this specification is a nuisance when the “constant” value is only constant under some weaker test like `string=` or `equal`.

It's especially annoying because, in SBCL, `defconstant` takes effect not only at load time but also at compile time, so that just compiling and loading reasonable code like

```
(defconstant +foobyte+ '(1 4))
```

runs into this undefined behavior. Many implementations of Common Lisp try to help the programmer around this annoyance by silently accepting the undefined code and trying to do what the programmer probably meant.

SBCL instead treats the undefined behavior as an error. Often such code can be rewritten in portable ANSI Common Lisp which has the desired behavior. E.g., the code above can be given an exactly defined meaning by replacing `defconstant` either with `defparameter` or with a customized macro which does the right thing, e.g.

```
(defmacro define-constant (name value &optional doc)
  '(defconstant ,name (if (boundp ',name) (symbol-value ',name) ,value)
    ,@(when doc (list doc))))
```

or possibly along the lines of the `defconstant-eqx` macro used internally in the implementation of SBCL itself. In circumstances where this is not appropriate, the programmer can handle the condition type `sb-ext:defconstant-uneql`, and choose either the `continue` or `abort` restart as appropriate.

2.3.5 Style Warnings

SBCL gives style warnings about various kinds of perfectly legal code, e.g.

- multiple `defuns` of the same symbol in different units;
- special variables not named in the conventional `*foo*` style, and lexical variables unconventionally named in the `*foo*` style

This causes friction with people who point out that other ways of organizing code (especially avoiding the use of `defgeneric`) are just as aesthetically stylish. However, these warnings should be read not as “warning, bad aesthetics detected, you have no style” but “warning, this style keeps the compiler from understanding the code as well as you might like.” That is, unless the compiler warns about such conditions, there's no way for the compiler to warn about some programming errors which would otherwise be easy to overlook. (Related bug: The warning about multiple `defuns` is pointlessly annoying when you compile and then load a function containing `defun` wrapped in `eval-when`, and ideally should be suppressed in that case, but still isn't as of SBCL 0.7.6.)

2.4 Development Tools

2.4.1 Editor Integration

Though SBCL can be used running “bare”, the recommended mode of development is with an editor connected to SBCL, supporting not only basic lisp editing (paren-matching, etc), but providing among other features an integrated debugger, interactive compilation, and automated documentation lookup.

Currently *SLIME*¹ (Superior Lisp Interaction Mode for Emacs) together with Emacs is recommended for use with SBCL, though other options exist as well.

SLIME can be downloaded from <http://www.common-lisp.net/project/slime/>.

2.4.2 Language Reference

CLHS (Common Lisp Hyperspec) is a hypertext version of the ANSI standard, made freely available by *LispWorks* – an invaluable reference.

See: <http://www.lispworks.com/reference/HyperSpec/index.html>

2.4.3 Generating Executables

SBCL can generate stand-alone executables. The generated executables include the SBCL runtime itself, so no restrictions are placed on program functionality. For example, a deployed program can call `compile` and `load`, which requires the compiler to be present in the executable. For further information, See [Function `sb-ext save-lisp-and-die`], page 10.

2.5 More SBCL Information

2.5.1 SBCL Homepage

The SBCL website at <http://www.sbcl.org/> has some general information, plus links to mailing lists devoted to SBCL, and to archives of these mailing lists. Subscribing to the mailing lists *sbcl-help* and *sbcl-announce* is recommended: both are fairly low-volume, and help you keep abreast with SBCL development.

2.5.2 Online Documentation

Documentation for non-ANSI extensions for various commands is available online from the SBCL executable itself. The extensions for functions which have their own command prompts (e.g. the debugger, and `inspect`) are documented in text available by typing `help` at their command prompts. The extensions for functions which don't have their own command prompt (such as `trace`) are described in their documentation strings, unless your SBCL was compiled with an option not to include documentation strings, in which case the documentation strings are only readable in the source code.

2.5.3 Additional Documentation Files

Besides this user manual both SBCL source and binary distributions include some other SBCL-specific documentation files, which should be installed along with this manual on your system, e.g. in `/usr/local/share/doc/sbcl/`.

COPYING	Licence and copyright summary.
CREDITS	Authorship information on various parts of SBCL.
INSTALL	Covers installing SBCL from both source and binary distributions on your system, and also has some installation related troubleshooting information.
NEWS	Summarizes changes between various SBCL versions.

2.5.4 Internals Documentation

If you're interested in the development of the SBCL system itself, then subscribing to *sbcl-devel* is a good idea.

SBCL internals documentation – besides comments in the source – is currently maintained as a *wiki-like* website: <http://sbcl-internals.cliki.net/>.

¹ Historically, the ILISP package at <http://ilisp.cons.org/> provided similar functionality, but it does not support modern SBCL versions.

Some low-level information describing the programming details of the conversion from CMUCL to SBCL is available in the `doc/FOR-CMUCL-DEVELOPERS` file in the SBCL distribution, though it is not installed by default.

2.6 More Common Lisp Information

2.6.1 Internet Community

The Common Lisp internet community is fairly diverse: `news://comp.lang.lisp` is fairly high volume newsgroup, but has a rather poor signal/noise ratio. Various special interest mailing lists and IRC tend to provide more content and less flames. `http://www.lisp.org` and `http://www.cliki.net` contain numerous pointers places in the net where lispers talk shop.

2.6.2 Third-party Libraries

For a wealth of information about free Common Lisp libraries and tools we recommend checking out *CLiki*: `http://www.cliki.net/`.

2.6.3 Common Lisp Books

If you're not a programmer and you're trying to learn, many introductory Lisp books are available. However, we don't have any standout favorites. If you can't decide, try checking the Usenet `news://comp.lang.lisp` FAQ for recent recommendations.

If you are an experienced programmer in other languages but need to learn about Common Lisp, some books stand out:

Practical Common Lisp, by Peter Seibel

An excellent introduction to the language, covering both the basics and "advanced topics" like macros, CLOS, and packages. Available both in print format and on the web: `http://www.gigamonkeys.com/book/`.

Paradigms Of Artificial Intelligence Programming, by Peter Norvig

Good information on general Common Lisp programming, and many nontrivial examples. Whether or not your work is AI, it's a very good book to look at.

On Lisp, by Paul Graham

An in-depth treatment of macros, but not recommended as a first Common Lisp book, since it is slightly pre-ANSI so you need to be on your guard against non-standard usages, and since it doesn't really even try to cover the language as a whole, focusing solely on macros. Downloadable from `http://www.paulgraham.com/onlisp.html`.

Object-Oriented Programming In Common Lisp, by Sonya Keene

With the exception of *Practical Common Lisp* most introductory books don't emphasize CLOS. This one does. Even if you're very knowledgeable about object oriented programming in the abstract, it's worth looking at this book if you want to do any OO in Common Lisp. Some abstractions in CLOS (especially multiple dispatch) go beyond anything you'll see in most OO systems, and there are a number of lesser differences as well. This book tends to help with the culture shock.

Art Of Metaobject Programming, by Gregor Kiczales et al.

Currently the prime source of information on the Common Lisp Metaobject Protocol, which is supported by SBCL. Section 2 (Chapters 5 and 6) are freely available at `http://mop.lisp.se/www.alu.org/mop/`.

2.7 History and Implementation of SBCL

You can work productively with SBCL without knowing or understanding anything about where it came from, how it is implemented, or how it extends the ANSI Common Lisp standard. However, a little knowledge can be helpful in order to understand error messages, to troubleshoot problems, to understand why some parts of the system are better debugged than others, and to anticipate which known bugs, known performance problems, and missing extensions are likely to be fixed, tuned, or added.

SBCL is descended from CMUCL, which is itself descended from Spice Lisp, including early implementations for the Mach operating system on the IBM RT, back in the 1980s. Some design decisions from that time are still reflected in the current implementation:

- The system expects to be loaded into a fixed-at-compile-time location in virtual memory, and also expects the location of all of its heap storage to be specified at compile time.
- The system overcommits memory, allocating large amounts of address space from the system (often more than the amount of virtual memory available) and then failing if ends up using too much of the allocated storage.
- The system is implemented as a C program which is responsible for supplying low-level services and loading a Lisp `.core` file.

SBCL also inherited some newer architectural features from CMUCL. The most important is that on some architectures it has a generational garbage collector (“GC”), which has various implications (mostly good) for performance. These are discussed in another chapter, Chapter 6 [Efficiency], page 37.

SBCL has diverged from CMUCL in that SBCL is now essentially a “compiler-only implementation” of Common Lisp. This is a change in implementation strategy, taking advantage of the freedom “any of these facilities might share the same execution strategy” guaranteed in the ANSI specification section 3.1 (“Evaluation”). It does not mean SBCL can’t be used interactively, and in fact the change is largely invisible to the casual user, since SBCL still can and does execute code interactively by compiling it on the fly. (It is visible if you know how to look, like using `compiled-function-p`; and it is visible in the way that SBCL doesn’t have many bugs which behave differently in interpreted code than in compiled code.) What it means is that in SBCL, the `eval` function only truly “interprets” a few easy kinds of forms, such as symbols which are `boundp`. More complicated forms are evaluated by calling `compile` and then calling `funcall` on the returned result.

The direct ancestor of SBCL is the x86 port of CMUCL. This port was in some ways the most cobbled-together of all the CMUCL ports, since a number of strange changes had to be made to support the register-poor x86 architecture. Some things (like tracing and debugging) do not work particularly well there. SBCL should be able to improve in these areas (and has already improved in some other areas), but it takes a while.

On the x86 SBCL – like the x86 port of CMUCL – uses a *conservative* GC. This means that it doesn’t maintain a strict separation between tagged and untagged data, instead treating some untagged data (e.g. raw floating point numbers) as possibly-tagged data and so not collecting any Lisp objects that they point to. This has some negative consequences for average time efficiency (though possibly no worse than the negative consequences of trying to implement an exact GC on a processor architecture as register-poor as the X86) and also has potentially unlimited consequences for worst-case memory efficiency. In practice, conservative garbage collectors work reasonably well, not getting anywhere near the worst case. But they can occasionally cause odd patterns of memory usage.

The fork from CMUCL was based on a major rewrite of the system bootstrap process. CMUCL has for many years tolerated a very unusual “build” procedure which doesn’t actually build the complete system from scratch, but instead progressively overwrites parts of a running system with new versions. This quasi-build procedure can cause various bizarre bootstrapping hangups, especially when a major change is made to the system. It also makes the connection between the current source code and the current executable more tenuous than in other software systems – it’s easy to accidentally “build” a CMUCL system containing characteristics not reflected in the current version of the source code.

Other major changes since the fork from CMUCL include

- SBCL has removed many CMUCL extensions, (e.g. IP networking, remote procedure call, Unix system interface, and X11 interface) from the core system. Most of these are available as contributed modules (distributed with SBCL) or third-party modules instead.
- SBCL has deleted or deprecated some nonstandard features and code complexity which helped efficiency at the price of maintainability. For example, the SBCL compiler no longer implements memory pooling internally (and so is simpler and more maintainable, but generates more garbage and runs more slowly), and various block-compilation efficiency-increasing extensions to the language have been deleted or are no longer used in the implementation of SBCL itself.

3 Starting and Stopping

3.1 Starting SBCL

3.1.1 From Shell to Lisp

To run SBCL type `sbcl` at the command line.

You should end up in the toplevel *REPL* (read, eval, print -loop), where you can interact with SBCL by typing expressions.

```
$ sbcl
This is SBCL 0.8.13.60, an implementation of ANSI Common Lisp.
More information about SBCL is available at <http://www.sbcl.org/>.

SBCL is free software, provided as is, with absolutely no warranty.
It is mostly in the public domain; some portions are provided under
BSD-style licenses. See the CREDITS and COPYING files in the
distribution for more information.
* (+ 2 2)

4
* (exit)
$
```

See also Section 3.3 [Command Line Options], page 12, and Section 3.2 [Stopping SBCL], page 9.

3.1.2 Running from Emacs

To run SBCL as an inferior-lisp from Emacs in your `.emacs` do something like:

```
;;; The SBCL binary and command-line arguments
(setq inferior-lisp-program "/usr/local/bin/sbcl --noinform")
```

For more information on using SBCL with Emacs, see Section 2.4.1 [Editor Integration], page 6.

3.1.3 Shebang Scripts

Standard Unix tools that are interpreters follow a common command line protocol that is necessary to work with “shebang scripts”. SBCL supports this via the `--script` command line option.

Example file (`hello.lisp`):

```
#!/usr/local/bin/sbcl --script
(write-line "Hello, World!")
```

Usage examples:

```
$ ./hello.lisp
Hello, World!
$ sbcl --script hello.lisp
Hello, World!
```

3.2 Stopping SBCL

3.2.1 Exit

SBCL can be stopped at any time by calling `sb-ext:exit`, optionally returning a specified numeric value to the calling process. See Chapter 13 [Threading], page 96, for information about terminating individual threads.

`sb-ext:exit &key code abort timeout` [Function]

Terminates the process, causing `sbcl` to exit with `code`. `code` defaults to 0 when `abort` is false, and 1 when it is true.

When `abort` is false (the default), current thread is first unwound, `*exit-hooks*` are run, other threads are terminated, and standard output streams are flushed before `sbcl` calls `exit(3)` -- at which point `atexit(3)` functions will run. If multiple threads call `exit` with `abort` being false, the first one to call it will complete the protocol.

When `abort` is true, `sbcl` exits immediately by calling `_exit(2)` without unwinding stack, or calling exit hooks. Note that `_exit(2)` does not call `atexit(3)` functions unlike `exit(3)`.

Recursive calls to `exit` cause `exit` to behave as if `abort` was true.

`timeout` controls waiting for other threads to terminate when `abort` is `nil`. Once current thread has been unwound and `*exit-hooks*` have been run, spawning new threads is prevented and all other threads are terminated by calling `terminate-thread` on them. The system then waits for them to finish using `join-thread`, waiting at most a total `timeout` seconds for all threads to join. Those threads that do not finish in time are simply ignored while the exit protocol continues. `timeout` defaults to `*exit-timeout*`, which in turn defaults to 60. `timeout nil` means to wait indefinitely.

Note that `timeout` applies only to `join-thread`, not `*exit-hooks*`. Since `terminate-thread` is asynchronous, getting multithreaded application termination with complex cleanups right using it can be tricky. To perform an orderly synchronous shutdown use an exit hook instead of relying on implicit thread termination.

Consequences are unspecified if serious conditions occur during `exit` excepting errors from `*exit-hooks*`, which cause warnings and stop execution of the hook that signaled, but otherwise allow the exit process to continue normally.

3.2.2 End of File

By default SBCL also exits on end of input, caused either by user pressing *Control-D* on an attached terminal, or end of input when using SBCL as part of a shell pipeline.

3.2.3 Saving a Core Image

SBCL has the ability to save its state as a file for later execution. This functionality is important for its bootstrapping process, and is also provided as an extension to the user.

`sb-ext:save-lisp-and-die` *core-file-name* &key *toplevel executable* [Function]
save-runtime-options callable-exports purify root-structures environment-name
compression

Save a "core image", i.e. enough information to restart a Lisp process later in the same state, in the file of the specified name. Only global state is preserved: the stack is unwound in the process.

The following &key arguments are defined:

:toplevel The function to run when the created core file is resumed. The default function handles command line `toplevel` option processing and runs the top level read-eval-print loop. This function returning is equivalent to `(sb-ext:exit :code 0)` being called.

`toplevel` functions should always provide an `abort` restart: otherwise code they call will run without one.

:executable

If true, arrange to combine the `sbcl` runtime and the core image to create a standalone executable. If false (the default), the core image will not be executable on its own. Executable images always behave as if they were passed the `-noinform` runtime option.

:save-runtime-options

If true, values of runtime options `-dynamic-space-size` and `-control-stack-size` that were used to start `sbcl` are stored in the standalone executable, and restored when the executable is run. This also inhibits normal runtime option processing, causing all command line arguments to be passed to the `toplevel`. Meaningless if `:executable` is `nil`.

:callable-exports

This should be a list of symbols to be initialized to the appropriate alien callables on startup. All exported symbols should be present as global symbols in the symbol table of the runtime before the saved core is loaded. When this list is non-empty, the `:toplevel` argument cannot be supplied.

:purify

If true (the default on `cheneygc`), do a purifying gc which moves all dynamically allocated objects into static space. This takes somewhat longer than the normal gc which is otherwise done, but it's only done once, and subsequent GC's will be done less often and will take less time in the resulting core file. See the `purify` function. This parameter has no effect on platforms using the generational garbage collector.

:root-structures

This should be a list of the main entry points in any newly loaded systems. This need not be supplied, but locality and/or gc performance may be better if they are. This has two different but related meanings: If **:purify** is true – and only for **cheneygc** – the root structures are those which anchor the set of objects moved into static space. On **gencgc** – and only on platforms supporting immobile code – these are the functions and/or function-names which commence a depth-first scan of code when reordering based on the statically observable call chain. The complete set of reachable objects is not affected per se. This argument is meaningless if neither enabling precondition holds.

:environment-name

This has no purpose; it is accepted only for legacy compatibility.

:compression

This is only meaningful if the runtime was built with the **:sb-core-compression** feature enabled. If **nil** (the default), saves to uncompressed core files. If **:sb-core-compression** was enabled at build-time, the argument may also be an integer from -1 to 9, corresponding to zlib compression levels, or **t** (which is equivalent to the default compression level, -1).

:application-type

Present only on Windows and is meaningful only with **:executable t**. Specifies the subsystem of the executable, **:console** or **:gui**. The notable difference is that **:gui** doesn't automatically create a console window. The default is **:console**.

The save/load process changes the values of some global variables:

standard-output, ***debug-io***, *etc.*

Everything related to open streams is necessarily changed, since the **os** won't let us preserve a stream across save and load.

default-pathname-defaults

This is reinitialized to reflect the working directory where the saved core is loaded.

save-lisp-and-die interacts with **sb-alien:load-shared-object**: see its documentation for details.

On threaded platforms only a single thread may remain running after **sb-ext:*save-hooks*** have run. Applications using multiple threads can be **save-lisp-and-die** friendly by registering a save-hook that quits any additional threads, and an init-hook that restarts them.

This implementation is not as polished and painless as you might like:

- It corrupts the current Lisp image enough that the current process needs to be killed afterwards. This can be worked around by forking another process that saves the core.
- There is absolutely no binary compatibility of core images between different runtime support programs. Even runtimes built from the same sources at different times are treated as incompatible for this purpose.

This isn't because we like it this way, but just because there don't seem to be good quick fixes for either limitation and no one has been sufficiently motivated to do lengthy fixes.

sb-ext:*save-hooks*

[Variable]

A list of function designators which are called in an unspecified order before creating a saved core image. Unused by **sbcl** itself: reserved for user and applications.

In cases where the standard initialization files have already been loaded into the saved core, and alternative ones should be used (or none at all), **SBCL** allows customizing the initfile pathname computation.

sb-ext:*sysinit-pathname-function*

[Variable]

Designator for a function of zero arguments called to obtain a pathname designator for the default sysinit file, or **nil**. If the function returns **nil**, no sysinit file is used unless one has been specified on the command-line.

sb-ext:*userinit-pathname-function*

[Variable]

Designator for a function of zero arguments called to obtain a pathname designator or a stream for the default userinit file, or **nil**. If the function returns **nil**, no userinit file is used unless one has been specified on the command-line.

To facilitate distribution of SBCL applications using external resources, the filesystem location of the SBCL core file being used is available from Lisp.

`sb-ext:*core-pathname*`

[Variable]

The absolute pathname of the running `sbcl` core.

3.2.4 Exit on Errors

SBCL can also be configured to exit if an unhandled error occurs, which is mainly useful for acting as part of a shell pipeline; doing so under most other circumstances would mean giving up large parts of the flexibility and robustness of Common Lisp. See Section 5.1 [Debugger Entry], page 26.

3.3 Command Line Options

Command line options can be considered an advanced topic; for ordinary interactive use, no command line arguments should be necessary.

In order to understand the command line argument syntax for SBCL, it is helpful to understand that the SBCL system is implemented as two components, a low-level runtime environment written in C and a higher-level system written in Common Lisp itself. Some command line arguments are processed during the initialization of the low-level runtime environment, some command line arguments are processed during the initialization of the Common Lisp system, and any remaining command line arguments are passed on to user code.

The full, unambiguous syntax for invoking SBCL at the command line is:

```
sbcl  runtime-option*  --end-runtime-options  toplevel-option*  --end-toplevel-options
user-options*
```

For convenience, the `--end-runtime-options` and `--end-toplevel-options` elements can be omitted. Omitting these elements can be convenient when you are running the program interactively, and you can see that no ambiguities are possible with the option values you are using. Omitting these elements is probably a bad idea for any batch file where any of the options are under user control, since it makes it impossible for SBCL to detect erroneous command line input, so that erroneous command line arguments will be passed on to the user program even if they were intended for the runtime system or the Lisp system.

3.3.1 Runtime Options

`--core` *corefilename*

Run the specified Lisp core file instead of the default. Note that if the Lisp core file is a user-created core file, it may run a nonstandard toplevel which does not recognize the standard toplevel options.

`--dynamic-space-size` *megabytes*

Size of the dynamic space reserved on startup in megabytes. Default value is platform dependent.

`--control-stack-size` *megabytes*

Size of control stack reserved for each thread in megabytes. Default value is 2.

`--noinform`

Suppress the printing of any banner or other informational message at startup. This makes it easier to write Lisp programs which work cleanly in Unix pipelines. See also the `--noprint` and `--disable-debugger` options.

`--disable-ldb`

Disable the low-level debugger. Only effective if SBCL is compiled with LDB.

`--lose-on-corruption`

There are some dangerous low level errors (for instance, control stack exhausted, memory fault) that (or whose handlers) can corrupt the image. By default SBCL prints a warning, then tries to continue and handle the error in Lisp, but this will not always work and SBCL may malfunction or even hang. With this option, upon encountering such an error SBCL will invoke `ldb` (if present and enabled) or else exit.

--script filename

As a runtime option this is equivalent to **--noinform --disable-ldb --lose-on-corruption --end-runtime-options --script filename**. See the description of **--script** as a toplevel option below. If there are no other command line arguments following **--script**, the filename argument can be omitted.

--merge-core-pages

When platform support is present, provide hints to the operating system that identical pages may be shared between processes until they are written to. This can be useful to reduce the memory usage on systems with multiple SBCL processes started from similar but differently-named core files, or from compressed cores. Without platform support, do nothing. By default only compressed cores trigger hinting.

--no-merge-core-pages

Ensures that no sharing hint is provided to the operating system.

--help Print some basic information about SBCL, then exit.**--version** Print SBCL's version information, then exit.

In the future, runtime options may be added to control behaviour such as lazy allocation of memory.

Runtime options, including any **--end-runtime-options** option, are stripped out of the command line before the Lisp toplevel logic gets a chance to see it.

3.3.2 Toplevel Options**--sysinit filename**

Load filename instead of the default system initialization file (see Section 3.4 [Initialization Files], page 14.)

--no-sysinit

Don't load a system-wide initialization file. If this option is given, the **--sysinit** option is ignored.

--userinit filename

Load filename instead of the default user initialization file (see Section 3.4 [Initialization Files], page 14.)

--no-userinit

Don't load a user initialization file. If this option is given, the **--userinit** option is ignored.

--eval command

After executing any initialization file, but before starting the read-eval-print loop on standard input, read and evaluate the command given. More than one **--eval** option can be used, and all will be read and executed, in the order they appear on the command line.

--load filename

This is equivalent to **--eval '(load "filename")'**. The special syntax is intended to reduce quoting headaches when invoking SBCL from shell scripts.

--noprint When ordinarily the toplevel "read-eval-print loop" would be executed, execute a "read-eval loop" instead, i.e. don't print a prompt and don't echo results. Combined with the **--noinform** runtime option, this makes it easier to write Lisp "scripts" which work cleanly in Unix pipelines.**--disable-debugger**

By default when SBCL encounters an error, it enters the builtin debugger, allowing interactive diagnosis and possible intercession. This option disables the debugger, causing errors to print a backtrace and exit with status 1 instead. When given, this option takes effect before loading of initialization files or processing **--eval** and **--load** options. See **sb-ext:disable-debugger** for details. See Section 5.1 [Debugger Entry], page 26.

--script filename

Implies **--no-userinit --no-sysinit --disable-debugger --end-toplevel-options**.

Causes the system to load the specified file instead of entering the read-eval-print-loop, and exit afterwards. If the file begins with a shebang line, it is ignored.

If there are no other command line arguments following, the filename can be omitted: this causes the script to be loaded from standard input instead. Shebang lines in standard input script are currently *not* ignored.

In either case, if there is an unhandled error (e.g. end of file, or a broken pipe) on either standard input, standard output, or standard error, the script silently exits with code 0. This allows e.g. safely piping output from SBCL to `head -n1` or similar.

3.4 Initialization Files

SBCL processes initialization files with `read` and `eval`, not `load`; hence initialization files can be used to set startup `*package*` and `*readtable*`, and for proclaiming a global optimization policy.

System Initialization File

Defaults to `$SBCL_HOME/sbclrc`, or if that doesn't exist to `/etc/sbclrc`. Can be overridden with the command line option `--sysinit` or `--no-sysinit` (see Section 3.3.2 [Toplevel Options], page 13).

The system initialization file is intended for system administrators and software packagers to configure locations of installed third party modules, etc.

User Initialization File

Defaults to `$HOME/.sbclrc`. Can be overridden with the command line option `--userinit` or `--no-userinit` (see Section 3.3.2 [Toplevel Options], page 13).

The user initialization file is intended for personal customizations, such as loading certain modules at startup, defining convenience functions to use in the REPL, handling automatic recompilation of FASLs (see Section 2.3.2 [FASL Format], page 4), etc.

Neither initialization file is required.

3.5 Initialization and Exit Hooks

SBCL provides hooks into the system initialization and exit.

`sb-ext:*init-hooks*`

[Variable]

A list of function designators which are called in an unspecified order when a saved core image starts up, after the system itself has been initialized, but before non-user threads such as the finalizer thread have been started.

Unused by `sbcl` itself: reserved for user and applications.

`sb-ext:*exit-hooks*`

[Variable]

A list of function designators which are called in an unspecified order when `sbcl` process exits.

Unused by `sbcl` itself: reserved for user and applications.

Using `(sb-ext:exit :abort t)`, or calling `exit(3)` directly circumvents these hooks.

4 Compiler

This chapter will discuss most compiler issues other than efficiency, including compiler error messages, the SBCL compiler's unusual approach to type safety in the presence of type declarations, the effects of various compiler optimization policies, and the way that inlining and open coding may cause optimized code to differ from a naive translation. Efficiency issues are sufficiently varied and separate that they have their own chapter, Chapter 6 [Efficiency], page 37.

4.1 Diagnostic Messages

4.1.1 Controlling Verbosity

The compiler can be quite verbose in its diagnostic reporting, rather more than some users would prefer – the amount of noise emitted can be controlled, however.

To control emission of compiler diagnostics (of any severity other than `error`: see Section 4.1.2 [Diagnostic Severity], page 16) use the `sb-ext:muffle-conditions` and `sb-ext:unmuffle-conditions` declarations, specifying the type of condition that is to be muffled (the muffling is done using an associated `muffle-warning` restart).

Global control:

```
;;; Muffle compiler-notes globally
(declare (sb-ext:muffle-conditions sb-ext:compiler-note))
```

Local control:

```
;;; Muffle compiler-notes based on lexical scope
(defun foo (x)
  (declare (optimize speed) (fixnum x)
           (sb-ext:muffle-conditions sb-ext:compiler-note))
  (values (* x 5) ; no compiler note from this
         (locally
          (declare (sb-ext:unmuffle-conditions sb-ext:compiler-note))
          ;; this one gives a compiler note
          (* x -5))))
```

`sb-ext:muffle-conditions`

[Declaration]

Syntax: `type*`

Muffles the diagnostic messages that would be caused by compile-time signals of given types.

`sb-ext:unmuffle-conditions`

[Declaration]

Syntax: `type*`

Cancels the effect of a previous `sb-ext:muffle-conditions` declaration.

Various details of *how* the compiler messages are printed can be controlled via the alist `sb-ext:*compiler-print-variable-alist*`.

`sb-ext:*compiler-print-variable-alist*`

[Variable]

an association list describing new bindings for special variables to be used by the compiler for error-reporting, etc. Eg.

```
(((*PRINT-LENGTH* . 10) (*PRINT-LEVEL* . 6) (*PRINT-PRETTY* . NIL))
```

The variables in the `car` positions are bound to the values in the `cdr` during the execution of some debug commands. When evaluating arbitrary expressions in the debugger, the normal values of the printer control variables are in effect.

Initially empty, `*compiler-print-variable-alist*` is typically used to specify bindings for printer control variables.

For information about muffling warnings signaled outside of the compiler, see Section 7.9 [Customization Hooks for Users], page 59.

4.1.2 Diagnostic Severity

There are four levels of compiler diagnostic severity:

1. error
2. warning
3. style warning
4. note

The first three levels correspond to condition classes which are defined in the ANSI standard for Common Lisp and which have special significance to the `compile` and `compile-file` functions. These levels of compiler error severity occur when the compiler handles conditions of these classes.

The fourth level of compiler error severity, *note*, corresponds to the `sb-ext:compiler-note`, and is used for problems which are too mild for the standard condition classes, typically hints about how efficiency might be improved. The `sb-ext:code-deletion-note`, a subtype of `compiler-note`, is signalled when the compiler deletes user-supplied code after proving that the code in question is unreachable.

Future work for SBCL includes expanding this hierarchy of types to allow more fine-grained control over emission of diagnostic messages.

`sb-ext:compiler-note`

[Condition]

Class precedence list: `compiler-note`, `condition`, `t`

Root of the hierarchy of conditions representing information discovered by the compiler that the user might wish to know, but which does not merit a `style-warning` (or any more serious condition).

`sb-ext:code-deletion-note`

[Condition]

Class precedence list: `code-deletion-note`, `compiler-note`, `condition`, `t`

A condition type signalled when the compiler deletes code that the user has written, having proved that it is unreachable.

4.1.3 Understanding Compile Diagnostics

The messages emitted by the compiler contain a lot of detail in a terse format, so they may be confusing at first. The messages will be illustrated using this example program:

```
(defmacro zoq (x)
  '(roq (ploq (+ ,x 3))))

(defun foo (y)
  (declare (symbol y))
  (zoq y))
```

The main problem with this program is that it is trying to add 3 to a symbol. Note also that the functions `roq` and `ploq` aren't defined anywhere.

4.1.3.1 The Parts of a Compiler Diagnostic

When processing this program, the compiler will produce this warning:

```
; file: /tmp/foo.lisp
; in: DEFUN FOO
;   (ZOQ Y)
; --> ROQ PLOQ
; ==>
;   (+ Y 3)
;
; caught WARNING:
;   Asserted type NUMBER conflicts with derived type (VALUES SYMBOL &OPTIONAL).
```

In this example we see each of the six possible parts of a compiler diagnostic:

1. `'file: /tmp/foo.lisp'` This is the name of the file that the compiler read the relevant code from. The file name is displayed because it may not be immediately obvious when there is an error during compilation of a large system, especially when `with-compilation-unit` is used to delay undefined warnings.

2. ‘in: DEFUN FOO’ This is the definition top level form responsible for the diagnostic. It is obtained by taking the first two elements of the enclosing form whose first element is a symbol beginning with “def”. If there is no such enclosing “def” form, then the outermost form is used. If there are multiple ‘def’ forms, then they are all printed from the outside in, separated by ‘=>’s. In this example, the problem was in the `defun` for `foo`.
3. ‘(ZOQ Y)’ This is the *original source* form responsible for the diagnostic. Original source means that the form directly appeared in the original input to the compiler, i.e. in the lambda passed to `compile` or in the top level form read from the source file. In this example, the expansion of the `zoq` macro was responsible for the message.
4. ‘--> ROQ PLOQ’ This is the *processing path* that the compiler used to produce the code that caused the message to be emitted. The processing path is a representation of the evaluated forms enclosing the actual source that the compiler encountered when processing the original source. The path is the first element of each form, or the form itself if the form is not a list. These forms result from the expansion of macros or source-to-source transformation done by the compiler. In this example, the enclosing evaluated forms are the calls to `roq` and `ploq`. These calls resulted from the expansion of the `zoq` macro.
5. ‘==> (+ Y 3)’ This is the *actual source* responsible for the diagnostic. If the actual source appears in the explanation, then we print the next enclosing evaluated form, instead of printing the actual source twice. (This is the form that would otherwise have been the last form of the processing path.) In this example, the problem is with the evaluation of the reference to the variable `y`.
6. ‘caught WARNING: Asserted type NUMBER conflicts with derived type (VALUES SYMBOL &OPTIONAL).’ This is the *explanation* of the problem. In this example, the problem is that, while the call to `+` requires that its arguments are all of type `number`, the compiler has derived that `y` will evaluate to a `symbol`. Note that ‘(VALUES SYMBOL &OPTIONAL)’ expresses that `y` evaluates to precisely one value.

Note that each part of the message is distinctively marked:

- ‘file:’ and ‘in:’ mark the file and definition, respectively.
- The original source is an indented form with no prefix.
- Each line of the processing path is prefixed with ‘-->’
- The actual source form is indented like the original source, but is marked by a preceding ‘==>’ line.
- The explanation is prefixed with the diagnostic severity, which can be ‘caught ERROR:’, ‘caught WARNING:’, ‘caught STYLE-WARNING:’, or ‘note:’.

Each part of the message is more specific than the preceding one. If consecutive messages are for nearby locations, then the front part of the messages would be the same. In this case, the compiler omits as much of the second message as in common with the first. For example:

```

; file: /tmp/foo.lisp
; in: DEFUN FOO
;   (ZOQ Y)
; --> ROQ
; ==>
;   (PLOQ (+ Y 3))
;
; caught STYLE-WARNING:
;   undefined function: PLOQ

; ==>
;   (ROQ (PLOQ (+ Y 3)))
;
; caught STYLE-WARNING:
;   undefined function: ROQ

```

In this example, the file, definition and original source are identical for the two messages, so the compiler omits them in the second message. If consecutive messages are entirely identical, then the compiler prints only the first message, followed by: ‘[Last message occurs *repeats* times]’ where *repeats* is the number of times the message was given.

If the source was not from a file, then no file line is printed. If the actual source is the same as the original source, then the processing path and actual source will be omitted. If no forms intervene between the original source and the actual source, then the processing path will also be omitted.

4.1.3.2 The Original and Actual Source

The *original source* displayed will almost always be a list. If the actual source for an message is a symbol, the original source will be the immediately enclosing evaluated list form. So even if the offending symbol does appear in the original source, the compiler will print the enclosing list and then print the symbol as the actual source (as though the symbol were introduced by a macro.)

When the *actual source* is displayed (and is not a symbol), it will always be code that resulted from the expansion of a macro or a source-to-source compiler optimization. This is code that did not appear in the original source program; it was introduced by the compiler.

Keep in mind that when the compiler displays a source form in an diagnostic message, it always displays the most specific (innermost) responsible form. For example, compiling this function

```
(defun bar (x)
  (let (a)
    (declare (fixnum a))
    (setq a (foo x))
    a))
```

gives this error message

```
; file: /tmp/foo.lisp
; in: DEFUN BAR
;   (LET (A)
;     (DECLARE (FIXNUM A))
;     (SETQ A (FOO X))
;     A)
;
; caught WARNING:
;   Asserted type FIXNUM conflicts with derived type (VALUES NULL &OPTIONAL).
```

This message is not saying “there is a problem somewhere in this `let`” – it is saying that there is a problem with the `let` itself. In this example, the problem is that `a`’s `nil` initial value is not a `fixnum`.

4.1.3.3 The Processing Path

The processing path is mainly useful for debugging macros, so if you don’t write macros, you can probably ignore it. Consider this example:

```
(defun foo (n)
  (dotimes (i n *undefined*)))
```

Compiling results in this error message:

```
; in: DEFUN FOO
;   (DOTIMES (I N *UNDEFINED*))
; --> DO BLOCK LET TAGBODY RETURN-FROM
; ==>
;   (PROGN *UNDEFINED*)
;
; caught WARNING:
;   undefined variable: *UNDEFINED*
```

Note that `do` appears in the processing path. This is because `dotimes` expands into:

```
(do ((i 0 (1+ i)) (#:g1 n))
    ((>= i #:g1) *undefined*)
    (declare (type unsigned-byte i)))
```

The rest of the processing path results from the expansion of `do`:

```
(block nil
  (let ((i 0) (#:g1 n))
    (declare (type unsigned-byte i))
    (tagbody (go #:g3)
      #:g2 (psetq i (1+ i))
```

```
#:g3      (unless (>= i #:g1) (go #:g2))
(return-from nil (progn *undefined*))))
```

In this example, the compiler descended into the `block`, `let`, `tagbody` and `return-from` to reach the `progn` printed as the actual source. This is a place where the “actual source appears in explanation” rule was applied. The innermost actual source form was the symbol `*undefined*` itself, but that also appeared in the explanation, so the compiler backed out one level.

4.2 Handling of Types

One of the most important features of the SBCL compiler (similar to the original CMUCL compiler, also known as *Python*) is its fairly sophisticated understanding of the Common Lisp type system and its conservative approach to the implementation of type declarations.

These two features reward the use of type declarations throughout development, even when high performance is not a concern. Also, as discussed in the chapter on performance (see Chapter 6 [Efficiency], page 37), the use of appropriate type declarations can be very important for performance as well.

The SBCL compiler also has a greater knowledge of the Common Lisp type system than other compilers. Support is incomplete only for types involving the `satisfies` type specifier.

4.2.1 Declarations as Assertions

The SBCL compiler treats type declarations differently from most other Lisp compilers. Under default compilation policy the compiler doesn’t blindly believe type declarations, but considers them assertions about the program that should be checked: all type declarations that have not been proven to always hold are asserted at runtime.

Remaining bugs in the compiler’s handling of types unfortunately provide some exceptions to this rule, see Section 4.2.4 [Implementation Limitations], page 21.

CLOS slot types form a notable exception. Types declared using the `:type` slot option in `defclass` are asserted if and only if the class was defined in *safe code* and the slot access location is in *safe code* as well. This laxness does not pose any internal consistency issues, as the CLOS slot types are not available for the type inferencer, nor do CLOS slot types provide any efficiency benefits.

There are three type checking policies available in SBCL, selectable via `optimize` declarations.

Full Type Checks

All declarations are considered assertions to be checked at runtime, and all type checks are precise. The default compilation policy provides full type checks.

Used when `(or (>= safety 2) (>= safety speed 1))`.

Weak Type Checks

Declared types may be simplified into faster to check supertypes: for example, `(or (integer -17 -7) (integer 7 17))` is simplified into `(integer -17 17)`.

Note: it is relatively easy to corrupt the heap when weak type checks are used if the program contains type-errors.

Used when `(and (< safety 2) (< safety speed))`

No Type Checks

All declarations are believed without assertions. Also disables argument count and array bounds checking.

Note: any type errors in code where type checks are not performed are liable to corrupt the heap.

Used when `(= safety 0)`.

4.2.2 Precise Type Checking

Precise checking means that the check is done as though `typep` had been called with the exact type specifier that appeared in the declaration.

If a variable is declared to be `(integer 3 17)` then its value must always be an integer between 3 and 17. If multiple type declarations apply to a single variable, then all the declarations must be correct; it is as though all the types were intersected producing a single `and` type specifier.

To gain maximum benefit from the compiler's type checking, you should always declare the types of function arguments and structure slots as precisely as possible. This often involves the use of `or`, `member`, and other list-style type specifiers.

4.2.3 Getting Existing Programs to Run

Since SBCL's compiler does much more comprehensive type checking than most Lisp compilers, SBCL may detect type errors in programs that have been debugged using other compilers. These errors are mostly incorrect declarations, although compile-time type errors can find actual bugs if parts of the program have never been tested.

Some incorrect declarations can only be detected by run-time type checking. It is very important to initially compile a program with full type checks (high `safety` optimization) and then test this safe version. After the checking version has been tested, then you can consider weakening or eliminating type checks. *This applies even to previously debugged programs*, because the SBCL compiler does much more type inference than other Common Lisp compilers, so an incorrect declaration can do more damage.

The most common problem is with variables whose constant initial value doesn't match the type declaration. Incorrect constant initial values will always be flagged by a compile-time type error, and they are simple to fix once located. Consider this code fragment:

```
(prog (foo)
  (declare (fixnum foo))
  (setq foo ...)
  ...)
```

Here `foo` is given an initial value of `nil`, but is declared to be a `fixnum`. Even if it is never read, the initial value of a variable must match the declared type. There are two ways to fix this problem. Change the declaration

```
(prog (foo)
  (declare (type (or fixnum null) foo))
  (setq foo ...)
  ...)
```

or change the initial value

```
(prog ((foo 0))
  (declare (fixnum foo))
  (setq foo ...)
  ...)
```

It is generally preferable to change to a legal initial value rather than to weaken the declaration, but sometimes it is simpler to weaken the declaration than to try to make an initial value of the appropriate type.

Another declaration problem occasionally encountered is incorrect declarations on `defmacro` arguments. This can happen when a function is converted into a macro. Consider this macro:

```
(defmacro my-1+ (x)
  (declare (fixnum x))
  `(the fixnum (1+ ,x)))
```

Although legal and well-defined Common Lisp code, this meaning of this definition is almost certainly not what the writer intended. For example, this call is illegal:

```
(my-1+ (+ 4 5))
```

This call is illegal because the argument to the macro is `(+ 4 5)`, which is a `list`, not a `fixnum`. Because of macro semantics, it is hardly ever useful to declare the types of macro arguments. If you really want to assert something about the type of the result of evaluating a macro argument, then put a `the` in the expansion:

```
(defmacro my-1+ (x)
  `(the fixnum (1+ (the fixnum ,x))))
```

In this case, it would be stylistically preferable to change this macro back to a function and declare it inline.

Some more subtle problems are caused by incorrect declarations that can't be detected at compile time. Consider this code:

```
(do ((pos 0 (position #\a string :start (1+ pos))))
  ((null pos))
```

```
(declare (fixnum pos))
...)
```

Although `pos` is almost always a `fixnum`, it is `nil` at the end of the loop. If this example is compiled with full type checks (the default), then running it will signal a type error at the end of the loop. If compiled without type checks, the program will go into an infinite loop (or perhaps `position` will complain because `(1+ nil)` isn't a sensible start.) Why? Because if you compile without type checks, the compiler just quietly believes the type declaration. Since the compiler believes that `pos` is always a `fixnum`, it believes that `pos` is never `nil`, so `(null pos)` is never true, and the loop exit test is optimized away. Such errors are sometimes flagged by unreachable code notes, but it is still important to initially compile and test any system with full type checks, even if the system works fine when compiled using other compilers.

In this case, the fix is to weaken the type declaration to `(or fixnum null)`¹.

Note that there is usually little performance penalty for weakening a declaration in this way. Any numeric operations in the body can still assume that the variable is a `fixnum`, since `nil` is not a legal numeric argument. Another possible fix would be to say:

```
(do ((pos 0 (position #\a string :start (1+ pos))))
  (null pos))
(let ((pos pos))
  (declare (fixnum pos))
  ...))
```

This would be preferable in some circumstances, since it would allow a non-standard representation to be used for the local `pos` variable in the loop body.

4.2.4 Implementation Limitations

Ideally, the compiler would consider *all* type declarations to be assertions, so that adding type declarations to a program, no matter how incorrect they might be, would *never* cause undefined behavior. However, the compiler is known to fall short of this goal in two areas:

- *Proclaimed* constraints on argument and result types of a function are supposed to be checked by the function. If the function type is proclaimed before function definition, type checks are inserted by the compiler, but the standard allows the reversed order, in which case the compiler will trust the declaration.
- The compiler cannot check types of an unknown number of values; if the number of generated values is unknown, but the number of consumed is known, only consumed values are checked.

For example,

```
(defun foo (x)
  (the integer (bar x)))
```

causes the following compiler diagnostic to be emitted:

```
; note: type assertion too complex to check:
; (VALUES INTEGER &REST T).
```

A partial workaround is instead write:

```
(defun foo (x)
  (the (values integer &optional) (bar x)))
```

These are important issues, but are not necessarily easy to fix, so they may, alas, remain in the system for a while.

4.3 Compiler Policy

Compiler policy is controlled by the `optimize` declaration, supporting all ANSI optimization qualities (`debug`, `safety`, `space`, and `speed`).²

For effects of various optimization qualities on type-safety and debuggability see Section 4.2.1 [Declarations as Assertions], page 19, and Section 5.6 [Debugger Policy Control], page 32.

¹ Actually, this declaration is unnecessary in SBCL, since it already knows that `position` returns a non-negative `fixnum` or `nil`.

² A deprecated extension `sb-ext:inhibit-warnings` is still supported, but liable to go away at any time.

Ordinarily, when the **speed** quality is high, the compiler emits notes to notify the programmer about its inability to apply various optimizations. For selective muffling of these notes See Section 4.1.1 [Controlling Verbosity], page 15.

The value of **space** mostly influences the compiler's decision whether to inline operations, which tend to increase the size of programs. Use the value 0 with caution, since it can cause the compiler to inline operations so indiscriminately that the net effect is to slow the program by causing cache misses or even swapping.

sb-ext:describe-compiler-policy *&optional spec* [Function]
Print all global optimization settings, augmented by **spec**.

sb-ext:restrict-compiler-policy *&optional quality min max* [Function]
Assign a minimum value to an optimization quality. **quality** is the name of the optimization quality to restrict, **min** (defaulting to zero) is the minimum allowed value, and **max** (defaults to 3) is the maximum. Returns the alist describing the current policy restrictions.
If **quality** is **nil** or not given, nothing is done.
Otherwise, if **min** is zero or **max** is 3 or neither are given, any existing restrictions of **quality** are removed. See also **:policy** option in **with-compilation-unit**.

cl:with-compilation-unit *options &body body* [Macro]
Affects compilations that take place within its dynamic extent. It is intended to be eg. wrapped around the compilation of all files in the same system.
Following options are defined:

:override *Boolean-Form*
One of the effects of this form is to delay undefined warnings until the end of the form, instead of giving them at the end of each compilation. If **override** is **nil** (the default), then the outermost **with-compilation-unit** form grabs the undefined warnings. Specifying **override** true causes that form to grab any enclosed warnings, even if it is enclosed by another **with-compilation-unit**.

:policy *Optimize-Declaration-Form*
Provides dynamic scoping for global compiler optimization qualities and restrictions, limiting effects of subsequent **optimize** proclamations and calls to **sb-ext:restrict-compiler-policy** to the dynamic scope of **body**.
If **override** is false, specified **policy** is merged with current global policy. If **override** is true, current global policy, including any restrictions, is discarded in favor of the specified **policy**.
Supplying **policy** **nil** is equivalent to the option not being supplied at all, ie. dynamic scoping of **policy** does not take place.
This option is an SBCL-specific experimental extension: Interface subject to change.

:source-namestring *Namestring-Form*
Attaches the value returned by the Namestring-Form to the internal debug-source information as the namestring of the source file. Normally the namestring of the input-file for **compile-file** is used: this option can be used to provide source-file information for functions compiled using **compile**, or to override the input-file of **compile-file**.
If both an outer and an inner **with-compilation-unit** provide a **source-namestring**, the inner one takes precedence. Unaffected by **:override**.
This is an SBCL-specific extension.

:source-plist *Plist-Form*
Attaches the value returned by the Plist-Form to internal debug-source information of functions compiled in within the dynamic extent of **body**.
Primarily for use by development environments, in order to eg. associate function definitions with editor-buffers. Can be accessed using **sb-introspect:definition-source-plist**.
If an outer **with-compilation-unit** form also provide a **source-plist**, it is appended to the end of the provided **source-plist**. Unaffected by **:override**.
This is an SBCL-specific extension.

Examples:

```
;; Prevent proclamations from the file leaking, and restrict
;; SAFETY to 3 -- otherwise uses the current global policy.
(with-compilation-unit (:policy '(optimize))
  (restrict-compiler-policy 'safety 3)
  (load "foo.lisp"))

;; Using default policy instead of the current global one,
;; except for DEBUG 3.
(with-compilation-unit (:policy '(optimize debug)
                        :override t)
  (load "foo.lisp"))

;; Same as if :POLICY had not been specified at all: SAFETY 3
;; proclamation leaks out from WITH-COMPILATION-UNIT.
(with-compilation-unit (:policy nil)
  (declare (optimize safety))
  (load "foo.lisp"))
```

4.4 Compiler Errors

4.4.1 Type Errors at Compile Time

If the compiler can prove at compile time that some portion of the program cannot be executed without a type error, then it will give a warning at compile time.

It is possible that the offending code would never actually be executed at run-time due to some higher level consistency constraint unknown to the compiler, so a type warning doesn't always indicate an incorrect program.

For example, consider this code fragment:

```
(defun raz (foo)
  (let ((x (case foo
             (:this 13)
             (:that 9)
             (:the-other 42))))
    (declare (fixnum x))
    (foo x)))
```

Compilation produces this warning:

```
; in: DEFUN RAZ
;   (CASE FOO (:THIS 13) (:THAT 9) (:THE-OTHER 42))
; --> LET COND IF COND IF COND IF
; ==>
;   (COND)
;
; caught WARNING:
;   This is not a FIXNUM:
;   NIL
```

In this case, the warning means that if `foo` isn't any of `:this`, `:that` or `:the-other`, then `x` will be initialized to `nil`, which the `fixnum` declaration makes illegal. The warning will go away if `ecase` is used instead of `case`, or if `:the-other` is changed to `t`.

This sort of spurious type warning happens moderately often in the expansion of complex macros and in inline functions. In such cases, there may be dead code that is impossible to correctly execute. The compiler can't always prove this code is dead (could never be executed), so it compiles the erroneous code (which will always signal an error if it is executed) and gives a warning.

4.4.2 Errors During Macroexpansion

The compiler handles errors that happen during macroexpansion, turning them into compiler errors. If you want to debug the error (to debug a macro), you can set `*break-on-signals*` to `error`. For example, this definition:

```
(defun foo (e l)
  (do ((current l (cdr current))
      ((atom current) nil))
    (when (eq (car current) e) (return current))))
```

gives this error:

```
; in: DEFUN FOO
;   (DO ((CURRENT L (CDR CURRENT))
;       ((ATOM CURRENT) NIL))
;     (WHEN (EQ (CAR CURRENT) E) (RETURN CURRENT)))
;
; caught ERROR:
;   (in macroexpansion of (DO # #))
;   (hint: For more precise location, try *BREAK-ON-SIGNALS*.)
;   DO step variable is not a symbol: (ATOM CURRENT)
```

4.4.3 Read Errors

SBCL's compiler does not attempt to recover from read errors when reading a source file, but instead just reports the offending character position and gives up on the entire source file.

4.5 Open Coding and Inline Expansion

Since Common Lisp forbids the redefinition of standard functions, the compiler can have special knowledge of these standard functions embedded in it. This special knowledge is used in various ways (open coding, inline expansion, source transformation), but the implications to the user are basically the same:

- Attempts to redefine standard functions may be frustrated, since the function may never be called. Although it is technically illegal to redefine standard functions, users sometimes want to implicitly redefine these functions when they are debugging using the `trace` macro. Special-casing of standard functions can be inhibited using the `notinline` declaration, but even then some phases of analysis such as type inferencing are applied by the compiler.
- The compiler can have multiple alternate implementations of standard functions that implement different trade-offs of speed, space and safety. This selection is based on the compiler policy, Section 4.3 [Compiler Policy], page 21.

When a function call is *open coded*, inline code whose effect is equivalent to the function call is substituted for that function call. When a function call is *closed coded*, it is usually left as is, although it might be turned into a call to a different function with different arguments. As an example, if `nthcdr` were to be open coded, then

```
(nthcdr 4 foobar)
might turn into
(cdr (cdr (cdr (cdr foobar))))
or even
```

```
(do ((i 0 (1+ i))
    (list foobar (cdr foobar)))
  ((= i 4) list))
```

If `nth` is closed coded, then

```
(nth x l)
might stay the same, or turn into something like
(car (nthcdr x l))
```

In general, open coding sacrifices space for speed, but some functions (such as `car`) are so simple that they are always open-coded. Even when not open-coded, a call to a standard function may be transformed into a different function call (as in the last example) or compiled as *static call*. Static function call uses a more efficient calling convention that forbids redefinition.

4.6 Interpreter

By default SBCL implements `eval` by calling the native code compiler.

SBCL also includes an interpreter for use in special cases where using the compiler is undesirable, for example due to compilation overhead. Unlike in some other Lisp implementations, in SBCL interpreted code is not safer or more debuggable than compiled code.

`sb-ext:*evaluator-mode*`

[Variable]

Toggle between different evaluator implementations. If set to `:compile`, an implementation of `eval` that calls the compiler will be used. If set to `:interpret`, an interpreter will be used.

4.7 Advanced Compiler Use and Efficiency Hints

For more advanced usages of the compiler, please see the chapter of the same name in the CMUCL manual. Many aspects of the compiler have stayed exactly the same, and there is a much more detailed explanation of the compiler's behavior and how to maximally optimize code in their manual. In particular, while SBCL no longer supports byte-code compilation, it does support CMUCL's block compilation facility allowing whole program optimization and increased use of the local call convention.

Unlike CMUCL, SBCL is able to open-code forward-referenced type tests while block compiling. This helps for mutually referential `defstructs` in particular.

5 Debugger

This chapter documents the debugging facilities of SBCL, including the debugger, single-stepper and `trace`, and the effect of `(optimize debug)` declarations.

5.1 Debugger Entry

5.1.1 Debugger Banner

When you enter the debugger, it looks something like this:

```
debugger invoked on a TYPE-ERROR in thread 11184:
  The value 3 is not of type LIST.
```

You can type `HELP` for debugger help, or `(SB-EXT:QUIT)` to exit from SBCL.

```
restarts (invokable by number or by possibly-abbreviated name):
  0: [ABORT  ] Reduce debugger level (leaving debugger, returning to toplevel).
  1: [TOPLEVEL] Restart at toplevel READ/EVAL/PRINT loop.
(CAR 1 3)
0]
```

The first group of lines describe what the error was that put us in the debugger. In this case `car` was called on `3`, causing a `type-error`.

This is followed by the “beginner help line”, which appears only if `sb-debug:*debug-beginner-help-p*` is true (default).

Next comes a listing of the active restart names, along with their descriptions – the ways we can restart execution after this error. In this case, both options return to top-level. Restarts can be selected by entering the corresponding number or name.

The current frame appears right underneath the restarts, immediately followed by the debugger prompt.

5.1.2 Debugger Invocation

The debugger is invoked when:

- `error` is called, and the condition it signals is not handled.
- `break` is called, or `signal` is called with a condition that matches the current `*break-on-signals*`.
- the debugger is explicitly entered with the `invoke-debugger` function.

When the debugger is invoked by a condition, ANSI mandates that the value of `*debugger-hook*`, if any, be called with two arguments: the condition that caused the debugger to be invoked and the previous value of `*debugger-hook*`. When this happens, `*debugger-hook*` is bound to `NIL` to prevent recursive errors. However, ANSI also mandates that `*debugger-hook*` not be invoked when the debugger is to be entered by the `break` function. For users who wish to provide an alternate debugger interface (and thus catch `break` entries into the debugger), SBCL provides `sb-ext:*invoke-debugger-hook*`, which is invoked during any entry into the debugger.

`sb-ext:*invoke-debugger-hook*` [Variable]

This is either `nil` or a designator for a function of two arguments, to be run when the debugger is about to be entered. The function is run with `*invoke-debugger-hook*` bound to `nil` to minimize recursive errors, and receives as arguments the condition that triggered debugger entry and the previous value of `*invoke-debugger-hook*`.

This mechanism is an `sbcl` extension similar to the standard `*debugger-hook*`. In contrast to `*debugger-hook*`, it is observed by `invoke-debugger` even when called by `break`.

5.2 Debugger Command Loop

The debugger is an interactive read-eval-print loop much like the normal top level, but some symbols are interpreted as debugger commands instead of being evaluated. A debugger command starts with the symbol name of the command, possibly followed by some arguments on the same line. Some commands prompt for

additional input. Debugger commands can be abbreviated by any unambiguous prefix: `help` can be typed as `h`, `he`, etc.

The package is not significant in debugger commands; any symbol with the name of a debugger command will work. If you want to show the value of a variable that happens also to be the name of a debugger command you can wrap the variable in a `progn` to hide it from the command loop.

The debugger prompt is “*frame*”, where *frame* is the number of the current frame. Frames are numbered starting from zero at the top (most recent call), increasing down to the bottom. The current frame is the frame that commands refer to.

It is possible to override the normal printing behaviour in the debugger by using the `sb-ext:*debug-print-variable-alist*`.

`sb-ext:*debug-print-variable-alist*` [Variable]

an association list describing new bindings for special variables to be used within the debugger. Eg.

((`*PRINT-LENGTH*` . 10) (`*PRINT-LEVEL*` . 6) (`*PRINT-PRETTY*` . NIL))

The variables in the `car` positions are bound to the values in the `cdr` during the execution of some debug commands. When evaluating arbitrary expressions in the debugger, the normal values of the printer control variables are in effect.

Initially empty, `*debug-print-variable-alist*` is typically used to provide bindings for printer control variables.

5.3 Stack Frames

A *stack frame* is the run-time representation of a call to a function; the frame stores the state that a function needs to remember what it is doing. Frames have:

- *variables* (see Section 5.4 [Variable Access], page 29), which are the values being operated on.
- *arguments* to the call (which are really just particularly interesting variables).
- a current source location (see Section 5.5 [Source Location Printing], page 31), which is the place in the program where the function was running when it stopped to call another function, or because of an interrupt or error.

5.3.1 Stack Motion

These commands move to a new stack frame and print the name of the function and the values of its arguments in the style of a Lisp function call:

`up` [Debugger Command]
Move up to the next higher frame. More recent function calls are considered to be higher on the stack.

`down` [Debugger Command]
Move down to the next lower frame.

`top` [Debugger Command]
Move to the highest frame, that is, the frame where the debugger was entered.

`bottom` [Debugger Command]
Move to the lowest frame.

`frame [n]` [Debugger Command]
Move to the frame with the specified number. Prompts for the number if not supplied. The frame with number 0 is the frame where the debugger was entered.

5.3.2 How Arguments are Printed

A frame is printed to look like a function call, but with the actual argument values in the argument positions. So the frame for this call in the source:

`(myfun (+ 3 4) 'a)`

would look like this:

`(MYFUN 7 A)`

All keyword and optional arguments are displayed with their actual values; if the corresponding argument was not supplied, the value will be the default. So this call:

`(subseq "foo" 1)`

would look like this:

```
(SUBSEQ "foo" 1 3)
```

And this call:

```
(string-upcase "test case")
```

would look like this:

```
(STRING-UPCASE "test case" :START 0 :END NIL)
```

The arguments to a function call are displayed by accessing the argument variables. Although those variables are initialized to the actual argument values, they can be set inside the function; in this case the new value will be displayed.

&rest arguments are handled somewhat differently. The value of the rest argument variable is displayed as the spread-out arguments to the call, so:

```
(format t "~A is a ~A." "This" 'test)
```

would look like this:

```
(FORMAT T "~A is a ~A." "This" 'TEST)
```

Rest arguments cause an exception to the normal display of keyword arguments in functions that have both **&rest** and **&key** arguments. In this case, the keyword argument variables are not displayed at all; the rest arg is displayed instead. So for these functions, only the keywords actually supplied will be shown, and the values displayed will be the argument values, not values of the (possibly modified) variables.

If the variable for an argument is never referenced by the function, it will be deleted. The variable value is then unavailable, so the debugger prints ‘#<unused-arg>’ instead of the value. Similarly, if for any of a number of reasons the value of the variable is unavailable or not known to be available (see Section 5.4 [Variable Access], page 29), then ‘#<unavailable-arg>’ will be printed instead of the argument value.

Note that inline expansion and open-coding affect what frames are present in the debugger, see Section 5.6 [Debugger Policy Control], page 32.

5.3.3 Function Names

If a function is defined by **defun** it will appear in backtrace by that name. Functions defined by **labels** and **flet** will appear as (**FLET name**) and (**LABELS name**) respectively. Anonymous lambdas will appear as (**LAMBDA lambda-list**).

5.3.3.1 Entry Point Details

Sometimes the compiler introduces new functions that are used to implement a user function, but are not directly specified in the source. This is mostly done for argument type and count checking.

With recursive or block compiled functions, an additional **external** frame may appear before the frame representing the first call to the recursive function or entry to the compiled block. This is a consequence of the way the compiler works: there is nothing odd with your program. You may also see **cleanup** frames during the execution of **unwind-protect** cleanup code, and **optional** for variable argument entry points.

5.3.4 Debug Tail Recursion

The compiler is “properly tail recursive.” If a function call is in a tail-recursive position, the stack frame will be deallocated *at the time of the call*, rather than after the call returns. Consider this backtrace:

```
(BAR ...)  
(FOO ...)
```

Because of tail recursion, it is not necessarily the case that **FOO** directly called **BAR**. It may be that **FOO** called some other function **FOO2** which then called **BAR** tail-recursively, as in this example:

```
(defun foo ()  
  ...  
  (foo2 ...)  
  ...)  
  
(defun foo2 (...)  
  ...  
  (bar ...))  
  
(defun bar (...)
```

...)

Usually the elimination of tail-recursive frames makes debugging more pleasant, since these frames are mostly uninformative. If there is any doubt about how one function called another, it can usually be eliminated by finding the source location in the calling frame. See Section 5.5 [Source Location Printing], page 31.

The elimination of tail-recursive frames can be prevented by disabling tail-recursion optimization, which happens when the `debug` optimization quality is greater than 2. See Section 5.6 [Debugger Policy Control], page 32.

5.3.5 Unknown Locations and Interrupts

The debugger operates using special debugging information attached to the compiled code. This debug information tells the debugger what it needs to know about the locations in the code where the debugger can be invoked. If the debugger somehow encounters a location not described in the debug information, then it is said to be *unknown*. If the code location for a frame is unknown, then some variables may be inaccessible, and the source location cannot be precisely displayed.

There are three reasons why a code location could be unknown:

- There is inadequate debug information due to the value of the `debug` optimization quality. See Section 5.6 [Debugger Policy Control], page 32.
- The debugger was entered because of an interrupt such as `C-c`.
- A hardware error such as “**bus error**” occurred in code that was compiled unsafely due to the value of the `safety` optimization quality.

In the last two cases, the values of argument variables are accessible, but may be incorrect. For more details on when variable values are accessible, see Section 5.4.1 [Variable Value Availability], page 30.

It is possible for an interrupt to happen when a function call or return is in progress. The debugger may then flame out with some obscure error or insist that the bottom of the stack has been reached, when the real problem is that the current stack frame can’t be located. If this happens, return from the interrupt and try again.

5.4 Variable Access

There are two ways to access the current frame’s local variables in the debugger: `list-locals` and `sb-debug:var`.

The debugger doesn’t really understand lexical scoping; it has just one namespace for all the variables in the current stack frame. If a symbol is the name of multiple variables in the same function, then the reference appears ambiguous, even though lexical scoping specifies which value is visible at any given source location. If the scopes of the two variables are not nested, then the debugger can resolve the ambiguity by observing that only one variable is accessible.

When there are ambiguous variables, the evaluator assigns each one a small integer identifier. The `sb-debug:var` function uses this identifier to distinguish between ambiguous variables. The `list-locals` command prints the identifier. In the following example, there are two variables named `X`. The first one has identifier 0 (which is not printed), the second one has identifier 1.

```
X = 1
X#1 = 2
```

`list-locals [prefix]` [Debugger Command]

This command prints the name and value of all variables in the current frame whose name has the specified *prefix*. *prefix* may be a string or a symbol. If no *prefix* is given, then all available variables are printed. If a variable has a potentially ambiguous name, then the name is printed with a “*#identifier*” suffix, where *identifier* is the small integer used to make the name unique.

`sb-debug:var name &optional identifier` [Function]

This function returns the value of the variable in the current frame with the specified *name*. If supplied, *identifier* determines which value to return when there are ambiguous variables.

When *name* is a symbol, it is interpreted as the symbol name of the variable, i.e. the package is significant. If *name* is an uninterned symbol (gensym), then return the value of the uninterned variable with the same name. If *name* is a string, `sb-debug:var` interprets it as the prefix of a variable name that must unambiguously complete to the name of a valid variable.

identifier is used to disambiguate the variable name; use `list-locals` to find out the identifiers.

5.4.1 Variable Value Availability

The value of a variable may be unavailable to the debugger in portions of the program where Lisp says that the variable is defined. If a variable value is not available, the debugger will not let you read or write that variable. With one exception, the debugger will never display an incorrect value for a variable. Rather than displaying incorrect values, the debugger tells you the value is unavailable.

The one exception is this: if you interrupt (e.g., with `C-c`) or if there is an unexpected hardware error such as “**bus error**” (which should only happen in unsafe code), then the values displayed for arguments to the interrupted frame might be incorrect.¹ This exception applies only to the interrupted frame: any frame farther down the stack will be fine.

The value of a variable may be unavailable for these reasons:

- The value of the `debug` optimization quality may have omitted debug information needed to determine whether the variable is available. Unless a variable is an argument, its value will only be available when `debug` is at least 2.
- The compiler did lifetime analysis and determined that the value was no longer needed, even though its scope had not been exited. Lifetime analysis is inhibited when the `debug` optimization quality is 3.
- The variable’s name is an uninterned symbol (gensym). To save space, the compiler only dumps debug information about uninterned variables when the `debug` optimization quality is 3.
- The frame’s location is unknown (see Section 5.3.5 [Unknown Locations and Interrupts], page 29) because the debugger was entered due to an interrupt or unexpected hardware error. Under these conditions the values of arguments will be available, but might be incorrect. This is the exception mentioned above.
- The variable (or the code referencing it) was optimized out of existence. Variables with no reads are always optimized away. The degree to which the compiler deletes variables will depend on the value of the `compilation-speed` optimization quality, but most source-level optimizations are done under all compilation policies.
- The variable is never set and its definition looks like

```
(LET ((var1 var2))
  ...)
```

In this case, `var1` is substituted with `var2`.

- The variable is never set and is referenced exactly once. In this case, the reference is substituted with the variable initial value.

Since it is especially useful to be able to get the arguments to a function, argument variables are treated specially when the `speed` optimization quality is less than 3 and the `debug` quality is at least 1. With this compilation policy, the values of argument variables are almost always available everywhere in the function, even at unknown locations. For non-argument variables, `debug` must be at least 2 for values to be available, and even then, values are only available at known locations.

5.4.2 Note On Lexical Variable Access

When the debugger command loop establishes variable bindings for available variables, these variable bindings have lexical scope and dynamic extent.² You can close over them, but such closures can’t be used as upward function arguments.

You can also set local variables using `setq`, but if the variable was closed over in the original source and never set, then setting the variable in the debugger may not change the value in all the functions the variable is defined in. Another risk of setting variables is that you may assign a value of a type that the compiler proved the variable could never take on. This may result in bad things happening.

¹ Since the location of an interrupt or hardware error will always be an unknown location, non-argument variable values will never be available in the interrupted frame. See Section 5.3.5 [Unknown Locations and Interrupts], page 29.

² The variable bindings are actually created using the Lisp `symbol-macrolet` special form.

5.5 Source Location Printing

One of the debugger's capabilities is source level debugging of compiled code. These commands display the source location for the current frame:

source [*context*] [Debugger Command]

This command displays the file that the current frame's function was defined from (if it was defined from a file), and then the source form responsible for generating the code that the current frame was executing. If *context* is specified, then it is an integer specifying the number of enclosing levels of list structure to print.

The source form for a location in the code is the innermost list present in the original source that encloses the form responsible for generating that code. If the actual source form is not a list, then some enclosing list will be printed. For example, if the source form was a reference to the variable `*some-random-special*`, then the innermost enclosing evaluated form will be printed. Here are some possible enclosing forms:

```
(let ((a *some-random-special*))
  ...)
```

```
(+ *some-random-special* ...)
```

If the code at a location was generated from the expansion of a macro or a source-level compiler optimization, then the form in the original source that expanded into that code will be printed. Suppose the file `/usr/me/mystuff.lisp` looked like this:

```
(defmacro mymac ()
  '(myfun))
```

```
(defun foo ()
  (mymac)
  ...)
```

If `foo` has called `myfun`, and is waiting for it to return, then the **source** command would print:

```
; File: /usr/me/mystuff.lisp
```

```
(MYMAC)
```

Note that the macro use was printed, not the actual function call form, `(myfun)`.

If enclosing source is printed by giving an argument to **source** or **vsources**, then the actual source form is marked by wrapping it in a list whose first element is `'#:***HERE***'`. In the previous example, **source** 1 would print:

```
; File: /usr/me/mystuff.lisp
```

```
(DEFUN FOO ()
  ( #:***HERE***
    (MYMAC))
  ...)
```

5.5.1 How the Source is Found

If the code was defined from Lisp by **compile** or **eval**, then the source can always be reliably located. If the code was defined from a `fasl` file created by **compile-file**, then the debugger gets the source forms it prints by reading them from the original source file. This is a potential problem, since the source file might have moved or changed since the time it was compiled.

The source file is opened using the **truename** of the source file pathname originally given to the compiler. This is an absolute pathname with all logical names and symbolic links expanded. If the file can't be located using this name, then the debugger gives up and signals an error.

If the source file can be found, but has been modified since the time it was compiled, the debugger prints this warning:

```
; File has been modified since compilation:
;   filename
; Using form offset instead of character position.
```

where *filename* is the name of the source file. It then proceeds using a robust but not foolproof heuristic for locating the source. This heuristic works if:

- No top-level forms before the top-level form containing the source have been added or deleted, and
- The top-level form containing the source has not been modified much. (More precisely, none of the list forms beginning before the source form have been added or deleted.)

If the heuristic doesn't work, the displayed source will be wrong, but will probably be near the actual source. If the “shape” of the top-level form in the source file is too different from the original form, then an error will be signaled. When the heuristic is used, the source location commands are noticeably slowed.

Source location printing can also be confused if (after the source was compiled) a read-macro you used in the code was redefined to expand into something different, or if a read-macro ever returns the same `eq` list twice. If you don't define read macros and don't use `##` in perverted ways, you don't need to worry about this.

5.5.2 Source Location Availability

Source location information is only available when the `debug` optimization quality is at least 2. If source location information is unavailable, the source commands will give an error message.

If source location information is available, but the source location is unknown because of an interrupt or unexpected hardware error (see Section 5.3.5 [Unknown Locations and Interrupts], page 29), then the command will print:

`Unknown location: using block start.`

and then proceed to print the source location for the start of the *basic block* enclosing the code location. It's a bit complicated to explain exactly what a basic block is, but here are some properties of the block start location:

- The block start location may be the same as the true location.
- The block start location will never be later in the program's flow of control than the true location.
- No conditional control structures (such as `if`, `cond`, or) will intervene between the block start and the true location (but note that some conditionals present in the original source could be optimized away.) Function calls *do not* end basic blocks.
- The head of a loop will be the start of a block.
- The programming language concept of “block structure” and the Lisp `block` special form are totally unrelated to the compiler's basic block.

In other words, the true location lies between the printed location and the next conditional (but watch out because the compiler may have changed the program on you.)

5.6 Debugger Policy Control

The compilation policy specified by `optimize` declarations affects the behavior seen in the debugger. The `debug` quality directly affects the debugger by controlling the amount of debugger information dumped. Other optimization qualities have indirect but observable effects due to changes in the way compilation is done.

Unlike the other optimization qualities (which are compared in relative value to evaluate tradeoffs), the `debug` optimization quality is directly translated to a level of debug information. This absolute interpretation allows the user to count on a particular amount of debug information being available even when the values of the other qualities are changed during compilation. These are the levels of debug information that correspond to the values of the `debug` quality:

- | | |
|-----|--|
| 0 | Only the function name and enough information to allow the stack to be parsed. |
| > 0 | Any level greater than 0 gives level 0 plus all argument variables. Values will only be accessible if the argument variable is never set and <code>speed</code> is not 3. SBCL allows any real value for optimization qualities. It may be useful to specify 0.5 to get backtrace argument display without argument documentation. |
| 1 | Level 1 provides argument documentation (printed argument lists) and derived argument/result type information. This makes <code>describe</code> more informative, and allows the compiler to do compile-time argument count and type checking for any calls compiled at run-time. This is the default. |

- 2 Level 1 plus all interned local variables, source location information, and lifetime information that tells the debugger when arguments are available (even when **speed** is 3 or the argument is set).
- > 2 Any level greater than 2 gives level 2 and in addition disables tail-call optimization, so that the backtrace will contain frames for all invoked functions, even those in tail positions.
- 3 Level 2 plus all uninterned variables. In addition, lifetime analysis is disabled (even when **speed** is 3), ensuring that all variable values are available at any known location within the scope of the binding. This has a speed penalty in addition to the obvious space penalty.
- > (max speed space)
 - If **debug** is greater than both **speed** and **space**, the command **return** can be used to continue execution by returning a value from the current stack frame.
- > (max speed space compilation-speed)
 - If **debug** is greater than all of **speed**, **space** and **compilation-speed** the code will be steppable (see Section 5.10 [Single Stepping], page 35).

As you can see, if the **speed** quality is 3, debugger performance is degraded. This effect comes from the elimination of argument variable special-casing (see Section 5.4.1 [Variable Value Availability], page 30). Some degree of speed/debuggability tradeoff is unavoidable, but the effect is not too drastic when **debug** is at least 2.

In addition to **inline** and **notinline** declarations, the relative values of the **speed** and **space** qualities also change whether functions are inline expanded. If a function is inline expanded, then there will be no frame to represent the call, and the arguments will be treated like any other local variable. Functions may also be “semi-inline”, in which case there is a frame to represent the call, but the call is to an optimized local version of the function, not to the original function.

5.7 Exiting Commands

These commands get you out of the debugger.

toplevel [Debugger Command]
Throw to top level.

restart *[n]* [Debugger Command]
Invokes the *n*th restart case as displayed by the **error** command. If *n* is not specified, the available restart cases are reported.

continue [Debugger Command]
Calls **continue** on the condition given to **debug**. If there is no restart case named *continue*, then an error is signaled.

abort [Debugger Command]
Calls **abort** on the condition given to **debug**. This is useful for popping debug command loop levels or aborting to top level, as the case may be.

return value [Debugger Command]
Returns *value* from the current stack frame. This command is available when the **debug** optimization quality is greater than both **speed** and **space**. Care must be taken that the value is of the same type as SBCL expects the stack frame to return.

restart-frame [Debugger Command]
Restarts execution of the current stack frame. This command is available when the **debug** optimization quality is greater than both **speed** and **space** and when the frame is for a global function. If the function is redefined in the debugger before the frame is restarted, the new function will be used.

5.8 Information Commands

Most of these commands print information about the current frame or function, but a few show general information.

help	[Debugger Command]
?	[Debugger Command]
Displays a synopsis of debugger commands.	
describe	[Debugger Command]
Calls describe on the current function and displays the number of local variables.	
print	[Debugger Command]
Displays the current function call as it would be displayed by moving to this frame.	
error	[Debugger Command]
Prints the condition given to invoke-debugger and the active proceed cases.	
backtrace [<i>n</i>]	[Debugger Command]
Displays all the frames from the current to the bottom. Only shows <i>n</i> frames if specified. The printing is controlled by <i>*debug-print-variable-alist*</i> .	

5.9 Function Tracing

The tracer causes selected functions to print their arguments and their results whenever they are called. Options allow conditional printing of the trace information and conditional breakpoints on function entry or exit.

c1:trace *&rest specs* [Macro]

trace {Option Global-Value}* {Name {Option Value}*}*
trace is a debugging tool that provides information when specified functions are called. In its simplest form:

(TRACE NAME-1 NAME-2 ...)

The NAMES are not evaluated. Each may be a symbol, denoting an individual function, or a string, denoting all functions fbound to symbols whose home package is the package with the given name.

Options allow modification of the default behavior. Each option is a pair of an option keyword and a value form. Global options are specified before the first name, and affect all functions traced by a given use of **trace**. Options may also be interspersed with function names, in which case they act as local options, only affecting tracing of the immediately preceding function name. Local options override global options.

By default, **trace** causes a printout on **trace-output** each time that one of the named functions is entered or returns. (This is the basic, **ansi** Common Lisp behavior of **trace**.)

The following options are defined:

:report *Report-Type*

If *Report-Type* is **trace** (the default) then information is reported by printing immediately. If *Report-Type* is **nil**, then the only effect of the trace is to execute other options (e.g. **print** or **break**).

:condition *Form*

:condition-after *Form*

:condition-all *Form*

If **:condition** is specified, then **trace** does nothing unless *Form* evaluates to true at the time of the call. **:condition-after** is similar, but suppresses the initial printout, and is tested when the function returns. **:condition-all** tries both before and after.

:break *Form*

:break-after *Form*

:break-all *Form*

If specified, and *Form* evaluates to true, then the debugger is invoked at the start of the function, at the end of the function, or both, according to the respective option.

:print *Form*

:print-after *Form*

:print-all *Form*

In addition to the usual printout, the result of evaluating *Form* is printed at the start of the function, at the end of the function, or both, according to the respective option. Multiple print options cause multiple values to be printed.

:wherein *Names*

If specified, *Names* is a function name or list of names. **trace** does nothing unless a call to one of those functions encloses the call to this function (i.e. it would appear in a backtrace.) Anonymous functions have string names like "DEFUN FOO".

:encapsulate *{:DEFAULT | t | NIL}*

If *t*, the default, tracing is done via encapsulation (redefining the function name) rather than by modifying the function. **:default** is not the default, but means to use encapsulation for interpreted functions and funcallable instances, breakpoints otherwise. When encapsulation is used, forms are **not** evaluated in the function's lexical environment, but **sb-debug:arg** can still be used.

:methods *{T | NIL}*

If *t*, any function argument naming a generic function will have its methods traced in addition to the generic function itself.

:function *Function-Form*

This is a not really an option, but rather another way of specifying what function to trace. The *Function-Form* is evaluated immediately, and the resulting function is traced.

:condition, **:break** and **:print** forms are evaluated in a context which mocks up the lexical environment of the called function, so that **sb-debug:var** and **sb-debug:arg** can be used. The **-after** and **-all** forms can use also use **sb-debug:arg**. In forms which are evaluated after the function call, (**sb-debug:arg** *n*) returns the *N*-th value returned by the function.

c1:untrace *&rest specs*

[Macro]

Remove tracing from the specified functions. Untraces all functions when called with no arguments.

sb-debug:*trace-indentation-step*

[Variable]

the increase in trace indentation at each call level

sb-debug:*max-trace-indentation*

[Variable]

If the trace indentation exceeds this value, then indentation restarts at 0.

sb-debug:*trace-encapsulate-default*

[Variable]

the default value for the **:encapsulate** option to **trace**

5.10 Single Stepping

SBCL includes an instrumentation based single-stepper for compiled code, that can be invoked via the **step** macro, or from within the debugger. See Section 5.6 [Debugger Policy Control], page 32, for details on enabling stepping for compiled code.

The following debugger commands are used for controlling single stepping.

start

[Debugger Command]

Selects the **continue** restart if one exists and starts single stepping. None of the other single stepping commands can be used before stepping has been started either by using **start** or by using the standard **step** macro.

step

[Debugger Command]

Steps into the current form. Stepping will be resumed when the next form that has been compiled with stepper instrumentation is evaluated.

next

[Debugger Command]

Steps over the current form. Stepping will be disabled until evaluation of the form is complete.

out

[Debugger Command]

Steps out of the current frame. Stepping will be disabled until the topmost stack frame that had been stepped into returns.

stop

[Debugger Command]

Stops the single stepper and resumes normal execution.

`cl:step` *form*

[Macro]

The form is evaluated with single stepping enabled. Function calls outside the lexical scope of the form can be stepped into only if the functions in question have been compiled with sufficient `debug` policy to be at least partially steppable.

5.11 Enabling and Disabling the Debugger

In certain contexts (e.g., non-interactive applications), it may be desirable to turn off the SBCL debugger (and possibly re-enable it). The functions here control the debugger.

`sb-ext:disable-debugger`

[Function]

When invoked, this function will turn off both the `sbcl` debugger and `ldb` (the low-level debugger). See also `enable-debugger`.

`sb-ext:enable-debugger`

[Function]

Restore the debugger if it has been turned off by `disable-debugger`.

6 Efficiency

6.1 Slot access

6.1.1 Structure object slot access

Structure slot accessors are efficient only if the compiler is able to open code them: compiling a call to a structure slot accessor before the structure is defined, declaring one `notinline`, or passing it as a functional argument to another function causes severe performance degradation.

6.1.2 Standard object slot access

The most efficient way to access a slot of a `standard-object` is by using `slot-value` with a constant slot name argument inside a `defmethod` body, where the variable holding the instance is a `specializer` parameter of the method and is never assigned to. The cost is roughly 1.6 times that of an open coded structure slot accessor.

Second most efficient way is to use a CLOS slot accessor, or `slot-value` with a constant slot name argument, but in circumstances other than specified above. This may be up to 3 times as slow as the method described above.

Example:

```
(defclass foo () ((bar)))

;; Fast: specializer and never assigned to
(defmethod quux ((foo foo) new)
  (let ((old (slot-value foo 'bar)))
    (setf (slot-value foo 'bar) new)
    old))

;; Slow: not a specializer
(defmethod quux ((foo foo) new)
  (let* ((temp foo)
        (old (slot-value temp 'bar)))
    (setf (slot-value temp 'bar) new)
    old))

;; Slow: assignment to FOO
(defmethod quux ((foo foo) new)
  (let ((old (slot-value foo 'bar)))
    (setf (slot-value foo 'bar) new)
    (setf foo new)
    old))
```

Note that when profiling code such as this, the first few calls to the generic function are not representative, as the dispatch mechanism is lazily set up during those calls.

6.2 Dynamic-extent allocation

SBCL has fairly extensive support for performing allocation on the stack when a variable is declared `dynamic-extent`. The `dynamic-extent` declarations are not verified, but are simply trusted as long as `sb-ext:*stack-allocate-dynamic-extent*` is true.

`sb-ext:*stack-allocate-dynamic-extent*` [Variable]

If true (the default), the compiler respects `dynamic-extent` declarations and stack allocates otherwise inaccessible parts of the object whenever possible. Potentially long (over one page in size) vectors are, however, not stack allocated except in zero `safety` code, as such a vector could overflow the stack without triggering overflow protection.

If dynamic extent constraints specified in the Common Lisp standard are violated, the best that can happen is for the program to have garbage in variables and return values; more commonly, the system will crash.

In particular, it is important to realize that dynamic extent is contagious:

```
(let* ((a (list 1 2 3))
      (b (cons a a)))
  (declare (dynamic-extent b))
  ;; Unless A is accessed elsewhere as well, SBCL will consider
  ;; it to be otherwise inaccessible -- it can only be accessed
  ;; through B, after all -- and stack allocate it as well.
  ;;
  ;; Hence returning (CAR B) here is unsafe.
  ...)
```

This allows stack allocation of complex structures. As a notable exception to this, SBCL does not as of 1.0.48.21 propagate dynamic-extentness through `&rest` arguments – but another conforming implementation might, so portable code should not rely on this.

```
(declare (inline foo))
(defun foo (fun &rest arguments)
  (declare (dynamic-extent arguments))
  (apply fun arguments))

(defun bar (a)
  ;; SBCL will heap allocate the result of (LIST A), and stack allocate
  ;; only the spine of the &rest list -- so this is safe, but unportable.
  ;;
  ;; Another implementation, including earlier versions of SBCL might consider
  ;; (LIST A) to be otherwise inaccessible and stack-allocate it as well!
  (foo #'car (list a)))
```

There are many cases when `dynamic-extent` declarations could be useful. At present, SBCL implements stack allocation for

- `&rest` lists, when these are declared `dynamic-extent`.
- `cons`, `list`, `list*`, and `vector` when the result is bound to a variable declared `dynamic-extent`.
- simple forms of `make-array`, whose result is bound to a variable declared `dynamic-extent`: stack allocation is possible only if the resulting array is known to be both simple and one-dimensional, and has a constant `:element-type`.

Note: stack space is limited, so allocation of a large vector may cause stack overflow. For this reason potentially large vectors, which might circumvent stack overflow detection, are stack allocated only in zero `safety` policies.

- closures defined with `flet` or `labels`, with a bound `dynamic-extent` declaration. Blocks and tags are also allocated on the heap, unless all non-local control transfers to them are compiled with zero `safety`.
- user-defined structures when the structure constructor defined using `defstruct` has been declared `inline` and the result of the call to the constructor is bound to a variable declared `dynamic-extent`.

Note: structures with “raw” slots can currently be stack-allocated only on x86 and x86-64. A “raw” slot is one whose declared type is a subtype of exactly one of: `double-float`, `single-float`, `(complex double-float)`, `(complex single-float)`, or `sb-ext:word`; but as an exception to the preceding, any subtype of `fixnum` is not stored as raw despite also being a subtype of `sb-ext:word`.

- all of the above when they appear as initial parts of another stack-allocated object.

Examples:

```
;;; Declaiming a structure constructor inline before definition makes
;;; stack allocation possible.
(declare (inline make-thing))
(defstruct thing obj next)

;;; Stack allocation of various objects bound to DYNAMIC-EXTENT
;;; variables.
(let* ((list (list 1 2 3))
      (nested (cons (list 1 2) (list* 3 4 (list 5)))))
  (vector (make-array 3 :element-type 'single-float)))
```

```

      (thing (make-thing :obj list
                        :next (make-thing :obj (make-array 3))))
      (declare (dynamic-extent list nested vector thing))
      ...)

;;; Stack allocation of arguments to a local function is equivalent
;;; to stack allocation of local variable values.
(flet ((f (x)
         (declare (dynamic-extent x))
         ...))
  (f (list 1 2 3))
  (f (cons (cons 1 2) (cons 3 4)))
  ...)

;;; Stack allocation of &REST lists
(defun foo (&rest args)
  (declare (dynamic-extent args))
  ...)

```

Future plans include

- Automatic detection of the common idiom of calling quantifiers with a closure, even when the closure is not declared `dynamic-extent`.

6.3 Modular arithmetic

Some numeric functions have a property: N lower bits of the result depend only on N lower bits of (all or some) arguments. If the compiler sees an expression of form `(logand exp mask)`, where *exp* is a tree of such “good” functions and *mask* is known to be of type `(unsigned-byte w)`, where w is a “good” width, all intermediate results will be cut to w bits (but it is not done for variables and constants!). This often results in an ability to use simple machine instructions for the functions.

Consider an example.

```

(defun i (x y)
  (declare (type (unsigned-byte 32) x y))
  (ldb (byte 32 0) (logxor x (lognot y))))

```

The result of `(lognot y)` will be negative and of type `(signed-byte 33)`, so a naive implementation on a 32-bit platform is unable to use 32-bit arithmetic here. But modular arithmetic optimizer is able to do it: because the result is cut down to 32 bits, the compiler will replace `logxor` and `lognot` with versions cutting results to 32 bits, and because terminals (here—expressions *x* and *y*) are also of type `(unsigned-byte 32)`, 32-bit machine arithmetic can be used.

As of SBCL 0.8.5 “good” functions are `+`, `-`; `logand`, `logior`, `logxor`, `lognot` and their combinations; and `ash` with the positive second argument. “Good” widths are 32 on 32-bit CPUs and 64 on 64-bit CPUs. While it is possible to support smaller widths as well, currently this is not implemented.

6.4 Global and Always-Bound variables

`sb-ext:defgroupal name value &optional doc` [Macro]

Defines *name* as a global variable that is always bound. *value* is evaluated and assigned to *name* both at compile- and load-time, but only if *name* is not already bound.

Global variables share their values between all threads, and cannot be locally bound, declared special, defined as constants, and neither bound nor defined as symbol macros.

See also the declarations `sb-ext:global` and `sb-ext:always-bound`.

`sb-ext:global`

[Declaration]

Syntax: `(sb-ext:global symbol*)`

Only valid as a global proclamation.

Specifies that the named symbols cannot be proclaimed or locally declared **special**. Proclaiming an already special or constant variable name as **global** signal an error. Allows more efficient value lookup in threaded environments in addition to expressing programmer intention.

sb-ext:always-bound

[Declaration]

Syntax: (**sb-ext:always-bound** symbol*)

Only valid as a global proclamation.

Specifies that the named symbols are always bound. Inhibits **makunbound** of the named symbols. Proclaiming an unbound symbol as **always-bound** signals an error. Allows the compiler to elide boundness checks from value lookups.

6.5 Miscellaneous Efficiency Issues

FIXME: The material in the CMUCL manual about getting good performance from the compiler should be reviewed, reformatted in Texinfo, lightly edited for SBCL, and substituted into this manual. In the meantime, the original CMUCL manual is still 95+% correct for the SBCL version of the Python compiler. See the sections

- Advanced Compiler Use and Efficiency Hints
- Advanced Compiler Introduction
- More About Types in Python
- Type Inference
- Source Optimization
- Tail Recursion
- Local Call
- Block Compilation
- Inline Expansion
- Object Representation
- Numbers
- General Efficiency Hints
- Efficiency Notes

Besides this information from the CMUCL manual, there are a few other points to keep in mind.

- The CMUCL manual doesn't seem to state it explicitly, but Python has a mental block about type inference when assignment is involved. Python is very aggressive and clever about inferring the types of values bound with **let**, **let***, inline function call, and so forth. However, it's much more passive and dumb about inferring the types of values assigned with **setq**, **setf**, and friends. It would be nice to fix this, but in the meantime don't expect that just because it's very smart about types in most respects it will be smart about types involved in assignments. (This doesn't affect its ability to benefit from explicit type declarations involving the assigned variables, only its ability to get by without explicit type declarations.)
- Since the time the CMUCL manual was written, CMUCL (and thus SBCL) has gotten a generational garbage collector. This means that there are some efficiency implications of various patterns of memory usage which aren't discussed in the CMUCL manual. (Some new material should be written about this.)
- SBCL has some important known efficiency problems. Perhaps the most important are
 - The garbage collector is not particularly efficient, at least on platforms without the generational collector (as of SBCL 0.8.9, all except x86).
 - Various aspects of the PCL implementation of CLOS are more inefficient than necessary.

Finally, note that Common Lisp defines many constructs which, in the infamous phrase, “could be compiled efficiently by a sufficiently smart compiler”. The phrase is infamous because making a compiler which actually is sufficiently smart to find all these optimizations systematically is well beyond the state of the art of current compiler technology. Instead, they're optimized on a case-by-case basis by hand-written code, or not optimized at all if the appropriate case hasn't been hand-coded. Some cases where no such hand-coding has been done as of SBCL version 0.6.3 include

- `(reduce #'f x)` where the type of `x` is known at compile time
- various bit vector operations, e.g. `(position 0 some-bit-vector)`
- specialized sequence idioms, e.g. `(remove item list :count 1)`
- cases where local compilation policy does not require excessive type checking, e.g. `(locally (declare (safety 1)) (assoc item list))` (which currently performs safe `endp` checking internal to `assoc`).

If your system's performance is suffering because of some construct which could in principle be compiled efficiently, but which the SBCL compiler can't in practice compile efficiently, consider writing a patch to the compiler and submitting it for inclusion in the main sources. Such code is often reasonably straightforward to write; search the sources for the string `"deftransform"` to find many examples (some straightforward, some less so).

7 Beyond the ANSI Standard

SBCL is derived from CMUCL, which implements many extensions to the ANSI standard. SBCL doesn't support as many extensions as CMUCL, but it still has quite a few. See Chapter 17 [Contributed Modules], page 119.

7.1 Reader Extensions

7.1.1 Extended Package Prefix Syntax

SBCL supports extended package prefix syntax, which allows specifying an alternate package instead of `*package*` for the reader to use as the default package for interning symbols:

package-name::form-with-interning-into-package

Example:

```
'foo::(bar quux zot) == '(foo::bar foo::quux foo::zot)
```

`*package*` is not rebound during the course of reading a form with extended package prefix syntax; if `foo::bar` would cause a read-time package lock violation, so does `foo::(bar)`.

7.1.2 Symbol Name Normalization

SBCL also extends the reader to normalize all symbols to Normalization Form KC in builds with Unicode enabled. Whether symbols are normalized is controlled by

`sb-ext:readtable-normalization readtable` [Function]

Returns `t` if `readtable` normalizes strings to `nfkc`, and `nil` otherwise. The `readtable-normalization` of the standard `readtable` is `t`.

Symbols created by `intern` and similar functions are not affected by this setting. If `sb-ext:readtable-normalization` is `t`, symbols that are not normalized are escaped during printing.

7.1.3 Decimal Syntax for Rationals

SBCL supports a decimal syntax for rationals, modelled after the standard syntax for floating-point numbers. If a number with floating-point syntax has an exponent marker of `r` or `R` (rather than one of the standard exponent markers), it is read as the rational with the exact value of the decimal number expressed as a float.

In addition, setting or binding the value of `*read-default-float-format*` to `rational` around a call to `read` or `read-from-string` has the effect that floating-point numbers without exponent markers are read as rational numbers, as if there had been an explicit `r` or `R` marker.

Floating point numbers of all types are printed with an exponent marker while the value of `*read-default-float-format*` is `rational`; however, rational numbers are printed in their standard syntax, irrespective of the value of `*read-default-float-format*`.

7.2 Package-Local Nicknames

SBCL allows giving packages local nicknames: they allow short and easy-to-use names to be used without fear of name conflict associated with normal nicknames.

A local nickname is valid only when inside the package for which it has been specified. Different packages can use same local nickname for different global names, or different local nickname for same global name.

Symbol `:package-local-nicknames` in `*features*` denotes the support for this feature.

`cl:defpackage name [[option]]* ⇒ package` [Macro]

Options are extended to include

- `:local-nicknames (local-nickname actual-package-name)*`

The package has the specified local nicknames for the corresponding actual packages.

Example:

```
(defpackage :bar (:intern "X"))
(defpackage :foo (:intern "X"))
(defpackage :quux (:use :cl) (:local-nicknames (:bar :foo) (:foo :bar)))
```

```
(find-symbol "X" :foo) ; => FOO::X
(find-symbol "X" :bar) ; => BAR::X
(let ((*package* (find-package :quux)))
  (find-symbol "X" :foo))           ; => BAR::X
(let ((*package* (find-package :quux)))
  (find-symbol "X" :bar))           ; => FOO::X
```

sb-ext:package-local-nicknames *package-designator* [Function]

Returns an alist of (local-nickname . actual-package) describing the nicknames local to the designated package.

When in the designated package, calls to **find-package** with the any of the local-nicknames will return the corresponding actual-package instead. This also affects all implied calls to **find-package**, including those performed by the reader.

When printing a package prefix for a symbol with a package local nickname, the local nickname is used instead of the real name in order to preserve print-read consistency.

See also: **add-package-local-nickname**, **package-locally-nicknamed-by-list**, **remove-package-local-nickname**, and the **defpackage** option **:local-nicknames**.

Experimental: interface subject to change.

sb-ext:package-locally-nicknamed-by-list *package-designator* [Function]

Returns a list of packages which have a local nickname for the designated package.

See also: **add-package-local-nickname**, **package-local-nicknames**, **remove-package-local-nickname**, and the **defpackage** option **:local-nicknames**.

Experimental: interface subject to change.

sb-ext:add-package-local-nickname *local-nickname actual-package &optional package-designator* [Function]

Adds **local-nickname** for **actual-package** in the designated package, defaulting to current package. **local-nickname** must be a string designator, and **actual-package** must be a package designator.

Returns the designated package.

Signals a continuable error if **local-nickname** is already a package local nickname for a different package, or if **local-nickname** is one of "CL", "COMMON-LISP", or, "KEYWORD", or if **local-nickname** is a global name or nickname for the package to which the nickname would be added.

When in the designated package, calls to **find-package** with the **local-nickname** will return the package the designated **actual-package** instead. This also affects all implied calls to **find-package**, including those performed by the reader.

When printing a package prefix for a symbol with a package local nickname, local nickname is used instead of the real name in order to preserve print-read consistency.

See also: **package-local-nicknames**, **package-locally-nicknamed-by-list**, **remove-package-local-nickname**, and the **defpackage** option **:local-nicknames**.

Experimental: interface subject to change.

sb-ext:remove-package-local-nickname *old-nickname &optional package-designator* [Function]

If the designated package had **old-nickname** as a local nickname for another package, it is removed. Returns true if the nickname existed and was removed, and **nil** otherwise.

See also: **add-package-local-nickname**, **package-local-nicknames**, **package-locally-nicknamed-by-list**, and the **defpackage** option **:local-nicknames**.

Experimental: interface subject to change.

7.3 Package Variance

Common Lisp standard specifies that “If the new definition is at variance with the current state of that package, the consequences are undefined;” SBCL by default signals a full warning and retains as much of the package state as possible.

This can be adjusted using **sb-ext:*on-package-variance***:

sb-ext:*on-package-variance*

[Variable]

Specifies behavior when redefining a package using **defpackage** and the definition is in variance with the current state of the package.

The value should be of the form:

```
(:WARN [T | packages-names] :ERROR [T | package-names])
```

specifying which packages get which behaviour -- with **t** signifying the default unless otherwise specified. If default is not specified, **:warn** is used.

:warn keeps as much state as possible and causes **sbcl** to signal a full warning.

:error causes **sbcl** to signal an error when the variant **defpackage** form is executed, with restarts provided for user to specify what action should be taken.

Example:

```
(setf *on-package-variance* '(:warn (:swank :swank-backend) :error t))
```

specifies to signal a warning if **swank** package is in variance, and an error otherwise.

7.4 Garbage Collection

SBCL provides additional garbage collection functionality not specified by ANSI.

sb-ext:*after-gc-hooks*

[Variable]

Called after each garbage collection, except for garbage collections triggered during thread exits. In a multithreaded environment these hooks may run in any thread.

sb-ext:gc &key full gen &allow-other-keys

[Function]

Initiate a garbage collection.

The default is to initiate a nursery collection, which may in turn trigger a collection of one or more older generations as well. If **full** is true, all generations are collected. If **gen** is provided, it can be used to specify the oldest generation guaranteed to be collected.

On CheneyGC platforms arguments **full** and **gen** take no effect: a full collection is always performed.

7.4.1 Finalization

Finalization allows code to be executed after an object has been garbage collected. This is useful for example for releasing foreign memory associated with a Lisp object.

sb-ext:finalize object function &key dont-save

[Function]

Arrange for the designated **function** to be called when there are no more references to **object**, including references in **function** itself.

If **dont-save** is true, the finalizer will be cancelled when **save-lisp-and-die** is called: this is useful for finalizers deallocating system memory, which might otherwise be called with addresses from the old image.

In a multithreaded environment **function** may be called in any thread. In both single and multithreaded environments **function** may be called in any dynamic scope: consequences are unspecified if **function** is not fully re-entrant.

Errors from **function** are handled and cause a **warning** to be signalled in whichever thread the **function** was called in.

Examples:

```
;;; GOOD, assuming RELEASE-HANDLE is re-entrant.
```

```
(let* ((handle (get-handle))
      (object (make-object handle)))
  (finalize object (lambda () (release-handle handle)))
  object)
```

```
;;; BAD, finalizer refers to object being finalized, causing
;;; it to be retained indefinitely!
```

```
(let* ((handle (get-handle))
      (object (make-object handle)))
  (finalize object
    (lambda ()
```

```

(release-handle (object-handle object))))))

;;; BAD, not re-entrant!
(defvar *rec* nil)

(defun oops ()
  (when *rec*
    (error "recursive OOPS")))
(let ((*rec* t))
  (gc))) ; or just cons enough to cause one

(progn
  (finalize "oops" #'oops)
  (oops)) ; GC causes re-entry to #'oops due to the finalizer
          ; -> ERROR, caught, WARNING signalled

sb-ext:cancel-finalization object [Function]
Cancel all finalizations for object.

```

7.4.2 Weak Pointers

Weak pointers allow references to objects to be maintained without keeping them from being garbage collected: useful for building caches among other things.

Hash tables can also have weak keys and values: see Section 7.12 [Hash Table Extensions], page 60.

```

sb-ext:make-weak-pointer object [Function]
Allocate and return a weak pointer which points to object.

sb-ext:weak-pointer-value weak-pointer [Function]
If weak-pointer is valid, return the value of weak-pointer and t. If the referent of weak-pointer has
been garbage collected, returns the values nil and nil.

```

7.4.3 Introspection and Tuning

```

sb-ext:*gc-run-time* [Variable]
Total cpu time spent doing garbage collection (as reported by get-internal-run-time.) Initialized to
zero on startup. It is safe to bind this to zero in order to measure gc time inside a certain section of
code, but doing so may interfere with results reported by eg. time.

sb-ext:bytes-consed-between-gcs [Function]
The amount of memory that will be allocated before the next garbage collection is initiated. This can be
set with setf.
On gencgc platforms this is the nursery size, and defaults to 5% of dynamic space size.
Note: currently changes to this value are lost when saving core.

sb-ext:dynamic-space-size [Function]
Size of the dynamic space in bytes.

sb-ext:get-bytes-consed [Function]
Return the number of bytes consed since the program began. Typically this result will be a consed bignum,
so if you have an application (e.g. profiling) which can't tolerate the overhead of consing bignums, you'll
probably want either to hack in at a lower level (as the code in the sb-profile package does), or to
design a more microefficient interface and submit it as a patch.

sb-ext:gc-logfile [Function]
Return the pathname used to log garbage collections. Can be setf. Default is nil, meaning collections
are not logged. If non-null, the designated file is opened before and after each collection, and generation
statistics are appended to it.

sb-ext:generation-average-age generation [Function]
Average age of memory allocated to generation: average number of times objects allocated to the
generation have seen younger objects promoted to it. Available on gencgc platforms only.
Experimental: interface subject to change.

```

sb-ext:generation-bytes-allocated *generation* [Function]
 Number of bytes allocated to **generation** currently. Available on **gencgc** platforms only.
 Experimental: interface subject to change.

sb-ext:generation-bytes-consed-between-gcs *generation* [Function]
 Number of bytes that can be allocated to **generation** before that generation is considered for garbage collection. This value is meaningless for generation 0 (the nursery): see **bytes-consed-between-gcs** instead. Default is 5% of the dynamic space size divided by the number of non-nursery generations. Can be assigned to using **setf**. Available on **gencgc** platforms only.
 Experimental: interface subject to change.

sb-ext:generation-minimum-age-before-gc *generation* [Function]
 Minimum average age of objects allocated to **generation** before that generation is may be garbage collected. Default is 0.75. See also **generation-average-age**. Can be assigned to using **setf**. Available on **gencgc** platforms only.
 Experimental: interface subject to change.

sb-ext:generation-number-of-gcs-before-promotion *generation* [Function]
 Number of times garbage collection is done on **generation** before automatic promotion to the next generation is triggered. Default is 1. Can be assigned to using **setf**. Available on **gencgc** platforms only.
 Experimental: interface subject to change.

sb-ext:generation-number-of-gcs *generation* [Function]
 Number of times garbage collection has been done on **generation** without promotion. Available on **gencgc** platforms only.
 Experimental: interface subject to change.

7.4.4 Tracing Live Objects Back to Roots

note: This feature is intended to help expert users diagnose rare low-level issues and should not be needed during normal usage. On top of that, the interface and implementation are experimental and may change at any time without further notice.

It is sometimes important to understand why a given object is retained in the Lisp image instead of being garbage collected. To help with this problem, SBCL provides a mechanism that searches through the different memory spaces, builds a path of references from a root to the object in question and finally reports this paths:

sb-ext:search-roots *weak-pointers &key criterion gc ignore print* [Function]
 Find roots keeping the targets of **weak-pointers** alive.

weak-pointers must be a single **sb-ext:weak-pointer** or a list of those, pointing to objects for which roots should be searched.

gc controls whether the search is performed in the context of a garbage collection, that is with all Lisp threads stopped. Possible values are:

t This is the more accurate of the object liveness proof generators, as there is no chance for other code to execute in between the garbage collection and production of the chain of referencing objects.

nil This works well enough, but might be adversely affected by actions of concurrent threads.

criterion determines just how rooty (how deep) a root must be in order to be considered. Possible values are:

:oldest This says we can stop upon seeing an object in the oldest gen to **gc**, or older. This is the easiest test to satisfy.

:pseudo-static This is usually the same as **:oldest**, unless the oldest gen to **gc** has been decreased.

:static To find a root of an image-backed object, you want to stop only at a truly **:static** object.

ignore is a list of objects to treat as if nonexistent in the heap. It can often be useful for finding a path to an interned symbol other than through its package by specifying the package as an ignored object.

print controls whether discovered paths should be returned or printed. Possible values are

true (other than :verbose)

nil Do not print any output. Instead return the discovered paths as a list of lists. Each list has the form

```
(target . (root node*))
```

where **target** is one of the target of one of the weak-pointers.

`root` is a description of the root at which the path starts and has one of the following forms:

```
:static If the root of the path is a non-collectible heap object.
```

:pinned If an unknown thread stack pins the root of the path.

((thread-name | thread-object) symbol currentp) If the path begins at a special binding of *symbol* in a thread. *currentp* is a boolean indicating whether the value is current or shadowed by another binding.

((thread-name | thread-object) guessed-pc) If the path begins at a lexical variable in the function whose code contains `guessed-pc`.

Each **node** in the remainder of the path is a cons (**object** . **slot**) indicating that the slot at index **slot** in **object** references the next path node.

Experimental: subject to change without prior notice.

An example of using this could look like this:

```
* (defvar *my-string* (list 1 2 "my string"))
*MY-STRING*
```

```
* (sb-ext:search-roots (sb-ext:make-weak-pointer (third *my-string*)))
```

```
-> ((SIMPLE-VECTOR 3)) #x10004E9EAF[2] -> (SYMBOL) #x5044100F[1] -> (CONS) #x100181FAE7[1] -> (
```

The single line of output on `cl:*standard-output*` shows the path from a root to "my string": the path starts with SBCL's internal package system data structures followed by the symbol (`cl-user:*my-string*`) followed the three cons cells of the list.

The `:print :verbose` argument produces similar behavior but describe the path elements in more detail:

```
* (sb-ext:search-roots (sb-ext:make-weak-pointer (third *my-string*)) :print :verbose)
```

Path to "my string":

```

6      10004E9EAF [ 2] a (simple-vector 3)
0      5044100F [ 1] COMMON-LISP-USER::*MY-STRING*
0      100181FAE7 [ 1] a cons
0      100181FAF7 [ 1] a cons
0      100181FB07 [ 0] a cons

```

The `:print nil` argument is a bit different:

```
* (sb-ext:search-roots (sb-ext:make-weak-pointer (third *my-string*))) :print nil)
```

```
(( "my string" :STATIC (#(*MY-STRING* 0 0) . 2) (*MY-STRING* . 1)
```

```
((1 2 "my string") . 1) ((2 "my string") . 1) (("my string") . 0)))
```

There is no output on `cl:*standard-output*` and the return value is a single path for the target object "my string". As before, the path shows the symbol and the three cons cells.

7.5 Metaobject Protocol

7.5.1 AMOP Compatibility of Metaobject Protocol

SBCL supports a metaobject protocol which is intended to be compatible with AMOP; present exceptions to this (as distinct from current bugs) are:

- `compute-effective-method` only returns one value, not two.

There is no record of what the second return value was meant to indicate, and apparently no clients for it.

- The direct superclasses of `funcallable-standard-object` are `(function standard-object)`, not `(standard-object function)`.

This is to ensure that the `standard-object` class is the last of the standardized classes before `t` appearing in the class precedence list of `generic-function` and `standard-generic-function`, as required by section 1.4.4.5 of the ANSI specification.

- the arguments `:declare` and `:declarations` to `ensure-generic-function` are both accepted, with the leftmost argument defining the declarations to be stored and returned by `generic-function-declarations`.

Where AMOP specifies `:declarations` as the keyword argument to `ensure-generic-function`, the Common Lisp standard specifies `:declare`. Portable code should use `:declare`.

- although SBCL obeys the requirement in AMOP that `validate-superclass` should treat `standard-class` and `funcallable-standard-class` as compatible metaclasses, we impose an additional requirement at class finalization time: a class of metaclass `funcallable-standard-class` must have `function` in its superclasses, and a class of metaclass `standard-class` must not.

After a class has been finalized, it is associated with a class prototype which is accessible by a standard mop function `class-prototype`. The user can then ask whether this object is a `function` or not in several different ways: whether it is a function according to `typep`; whether its `class-of` is `subtypep function`, or whether `function` appears in the superclasses of the class. The additional consistency requirement comes from the desire to make all of these answers the same.

The following class definitions are bad, and will lead to errors either immediately or if an instance is created:

```
(defclass bad-object (funcallable-standard-object)
  ()
  (:metaclass standard-class))

(defclass bad-funcallable-object (standard-object)
  ()
  (:metaclass funcallable-standard-class))
```

The following definition is acceptable:

```
(defclass mixin ()
  ((slot :initarg slot)))
(defclass funcallable-object (funcallable-standard-object mixin)
  ()
  (:metaclass funcallable-standard-class))
```

and leads to a class whose instances are funcallable and have one slot.

Note that this requirement also applies to the class `funcallable-standard-object`, which has metaclass `funcallable-standard-class` rather than `standard-class` as AMOP specifies.

- the requirement that “No portable class C_p may inherit, by virtue of being a direct or indirect subclass of a specified class, any slot for which the name is a symbol accessible in the `common-lisp-user` package or exported by any package defined in the ANSI Common Lisp standard.” is interpreted to mean that the standardized classes themselves should not have slots named by external symbols of public packages.

The rationale behind the restriction is likely to be similar to the ANSI Common Lisp restriction on defining functions, variables and types named by symbols in the Common Lisp package: preventing two independent pieces of software from colliding with each other.

- specializations of the `new-value` argument to `(setf slot-value-using-class)` are not allowed: all user-defined methods must have a `specializer` of the class `t`.

This prohibition is motivated by a separation of layers: the `slot-value-using-class` family of functions is intended for use in implementing different and new slot allocation strategies, rather than in performing application-level dispatching. Additionally, with this requirement, there is a one-to-one mapping between metaclass, class and slot-definition-class tuples and effective methods of `(setf slot-value-using-class)`, which permits optimization of `(setf slot-value-using-class)`’s discriminating function in the same manner as for `slot-value-using-class` and `slot-boundp-using-class`.

Note that application code may specialize on the `new-value` argument of slot accessors.

- the class named by the `name` argument to `ensure-class`, if any, is only redefined if it is the proper name of that class; otherwise, a new class is created.

This is consistent with the description of `ensure-class` in AMOP as the functional version of `defclass`, which has this behaviour; however, it is not consistent with the weaker requirement in AMOP, which states that any class found by `find-class`, no matter what its `class-name`, is redefined.

- an error is not signaled in the case of the `:name` initialization argument for `slot-definition` objects being a constant, when the slot definition is of type `structure-slot-definition` (i.e. it is associated with a class of type `structure-class`).

This allows code which uses constant names for structure slots to continue working as specified in ANSI, while enforcing the constraint for all other types of slot.

- the class named `t` is not an instance of the `built-in-class` metaclass.

AMOP specifies, in the “Inheritance Structure of Metaobject Classes” section, that the class named `t` should be an instance of `built-in-class`. However, it also specifies that `validate-superclass` should return true (indicating that a direct superclass relationship is permissible) if the second argument is the class named `t`. Also, ANSI specifies that classes with metaclass `built-in-class` may not be subclassed using `defclass`, and also that the class named `t` is the universal superclass, inconsistent with it being a `built-in-class`.

- uses of `change-class` and redefinitions of classes with `defclass` (or the functional interfaces `ensure-class` or `ensure-class-using-class`) must ensure that for each slot with allocation `:instance` or `:class`, the set of applicable methods on the `slot-value-using-class` family of generic functions is the same before and after the change.

This is required for correct operation of the protocol to update instances for the new or redefined class, and can be seen as part of the contract of the `:instance` or `:class` allocations.

7.5.2 Metaobject Protocol Extensions

In addition, SBCL supports extensions to the Metaobject protocol from AMOP; at present, they are:

- compile-time support for generating specializer metaobjects from specializer names in `defmethod` forms is provided by the `make-method-specializers-form` function, which returns a form which, when evaluated in the lexical environment of the `defmethod`, returns a list of specializer metaobjects. This operator suffers from similar restrictions to those affecting `make-method-lambda`, namely that the generic function must be defined when the `defmethod` form is expanded, so that the correct method of `make-method-specializers-form` is invoked. The system-provided method on `make-method-specializers-form` generates a call to `find-class` for each symbol specializer name, and a call to `intern-eql-specializer` for each `(eql x)` specializer name.
- run-time support for converting between specializer names and specializer metaobjects, mostly for the purposes of `find-method`, is provided by `parse-specializer-using-class` and `unparse-specializer-using-class`, which dispatch on their first argument, the generic function associated with a method with the given specializer. The system-provided methods on those methods convert between classes and proper names and between lists of the form `(eql x)` and interned eql specializer objects.
- distinguishing unbound instance allocated slots from bound ones when using `standard-instance-access` and `funcallable-standard-instance-access` is possible by comparison to the symbol-macro `+slot-unbound+`.

7.6 Extensible Sequences

ANSI Common Lisp has a class `sequence` with subclasses `list` and `vector` on which the “sequence functions” like `find`, `subseq`, etc. operate. As an extension to the ANSI specification, SBCL allows additional subclasses of `sequence` to be defined¹.

Users of this extension just make instances of `sequence` subclasses and transparently operate on them using sequence functions:

```
(coerce (subseq (make-instance 'my-sequence) 5 10) 'list)
```

¹ A motivation, rationale and additional examples for the design of this extension can be found in the paper *Rhodes, Christophe (2007): User-extensible sequences in Common Lisp* available for download at <http://www.doc.gold.ac.uk/~mas01cr/papers/ilc2007/sequences-20070301.pdf>.

From this perspective, no distinction between builtin and user-defined **sequence** subclasses should be necessary.

Providers of the extension, that is of user-defined **sequence** subclasses, have to adhere to a “sequence protocol” which consists of a set of generic functions in the **sequence** package.

A minimal **sequence** subclass has to specify **standard-object** and **sequence** as its superclasses and has to be the specialization of the **sequence** parameter of methods on at least the following generic functions:

sb-sequence:length *sequence* [Generic Function]
Returns the length of **sequence** or signals a **protocol-unimplemented** error if the sequence protocol is not implemented for the class of **sequence**.

sb-sequence:elt *sequence index* [Generic Function]
Returns the element at position **index** of **sequence** or signals a **protocol-unimplemented** error if the sequence protocol is not implemented for the class of **sequence**.

(**setf sb-sequence:elt**) [Generic Function]
Replaces the element at position **index** of **sequence** with **new-value** and returns **new-value** or signals a **protocol-unimplemented** error if the sequence protocol is not implemented for the class of **sequence**.

sb-sequence:adjust-sequence *sequence length &key initial-element* [Generic Function]
initial-contents
Return destructively modified **sequence** or a freshly allocated sequence of the same class as **sequence** of length **length**. Elements of the returned sequence are initialized to **initial-element**, if supplied, initialized to **initial-contents** if supplied, or identical to the elements of **sequence** if neither is supplied. Signals a **protocol-unimplemented** error if the sequence protocol is not implemented for the class of **sequence**.

sb-sequence:make-sequence-like *sequence length &key initial-element* [Generic Function]
initial-contents
Returns a freshly allocated sequence of length **length** and of the same class as **sequence**. Elements of the new sequence are initialized to **initial-element**, if supplied, initialized to **initial-contents** if supplied, or identical to the elements of **sequence** if neither is supplied. Signals a **protocol-unimplemented** error if the sequence protocol is not implemented for the class of **sequence**.

make-sequence-like is needed for functions returning freshly-allocated sequences such as **subseq** or **copy-seq**. **adjust-sequence** is needed for functions which destructively modify their arguments such as **delete**. In fact, all other sequence functions can be implemented in terms of the above functions and actually are, if no additional methods are defined. However, relying on these generic implementations, in particular not implementing the iterator protocol can incur a high performance penalty See Section 7.6.1 [Iterator Protocol], page 52.

When the sequence protocol is only partially implemented for a given **sequence** subclass, an attempt to apply one of the missing operations to instances of that class signals the following condition:

sb-sequence:protocol-unimplemented [Condition]
Class precedence list: **protocol-unimplemented**, **type-error**, **error**, **serious-condition**, **condition**, **t**
This error is signaled if a sequence operation is applied to an instance of a sequence class that does not support the operation.

In addition to the mandatory functions above, methods on the sequence functions listed below can be defined.

There are two noteworthy irregularities:

- The function **sb-sequence:empty** does not have a counterpart in the **cl** package. It is intended to be used instead of **length** when working with lazy or infinite sequences.
- The functions **map**, **concatenate** and **merge** receive a type designator specifying the type of the constructed sequence as their first argument. However, the corresponding generic functions **sb-sequence:map**, **sb-sequence:concatenate** and **sb-sequence:merge** receive a prototype instance of the requested **sequence** subclass instead.

sb-sequence:empty *sequence*

[Generic Function]

Returns **t** if **sequence** is an empty sequence and **nil** otherwise. Signals an error if **sequence** is not a sequence.

- **sb-sequence:count**, **sb-sequence:count-if**, **sb-sequence:count-if-not**
- **sb-sequence:find**, **sb-sequence:find-if**, **sb-sequence:find-if-not**
- **sb-sequence:position**, **sb-sequence:position-if**, **sb-sequence:position-if-not**
- **sb-sequence:subseq**
- **sb-sequence:copy-seq**
- **sb-sequence:fill**
-

sb-sequence:map *result-prototype function sequence &rest sequences*

[Generic Function]

Implements **cl:map** for extended sequences.

result-prototype corresponds to the **result-type** of **cl:map** but receives a prototype instance of an extended sequence class instead of a type specifier. By dispatching on **result-prototype**, methods on this generic function specify how extended sequence classes act when they are specified as the result type in a **cl:map** call. **result-prototype** may not be fully initialized and thus should only be used for dispatch and to determine its class.

Another difference to **cl:map** is that **function** is a function, not a function designator.

- **sb-sequence:nsubstitute**, **sb-sequence:nsubstitute-if**, **sb-sequence:nsubstitute-if-not**, **sb-sequence:substitute**, **sb-sequence:substitute-if**, **sb-sequence:substitute-if-not**
- **sb-sequence:replace**
- **sb-sequence:nreverse**, **sb-sequence:reverse**
-

sb-sequence:concatenate *result-prototype &rest sequences*

[Generic Function]

Implements **cl:concatenate** for extended sequences.

result-prototype corresponds to the **result-type** of **cl:concatenate** but receives a prototype instance of an extended sequence class instead of a type specifier. By dispatching on **result-prototype**, methods on this generic function specify how extended sequence classes act when they are specified as the result type in a **cl:concatenate** call. **result-prototype** may not be fully initialized and thus should only be used for dispatch and to determine its class.

- **sb-sequence:reduce**
- **sb-sequence:mismatch**
- **sb-sequence:search**
- **sb-sequence:delete**, **sb-sequence:delete-if**, **sb-sequence:delete-if-not**, **sb-sequence:remove**, **sb-sequence:remove-if**, **sb-sequence:remove-if-not**,
- **sb-sequence:delete-duplicates**, **sb-sequence:remove-duplicates**
- **sb-sequence:sort**, **sb-sequence:stable-sort**
-

sb-sequence:merge *result-prototype sequence1 sequence2 predicate &key key*

[Generic Function]

Implements **cl:merge** for extended sequences.

result-prototype corresponds to the **result-type** of **cl:merge** but receives a prototype instance of an extended sequence class instead of a type specifier. By dispatching on **result-prototype**, methods on this generic function specify how extended sequence classes act when they are specified as the result type in a **cl:merge** call. **result-prototype** may not be fully initialized and thus should only be used for dispatch and to determine its class.

Another difference to **cl:merge** is that **predicate** is a function, not a function designator.

In the spirit of **dolist**, generic sequences can be traversed using the macro

sb-sequence:dosequence (*element sequence &optional return*) *&body body*

[Macro]

Executes **body** with **element** subsequently bound to each element of **sequence**, then returns **return**.

7.6.1 Iterator Protocol

The iterator protocol allows subsequently accessing some or all elements of a sequence in forward or reverse direction. Users first call `make-sequence-iterator` to create an iteration state and receive functions to query and mutate it. These functions allow, among other things, moving to, retrieving or modifying elements of the sequence. An iteration state consists of a state object, a limit object, a from-end indicator and the following six functions to query or mutate this state:

step function *sequence iterator from-end* [Function]
 Moves the iterator one position forward or backward in the associated sequence depending on the iteration direction.

endp function *sequence iterator limit from-end* [Function]
 Returns non-`nil` when the iterator has reached the end of the associated sequence with respect to the iteration direction.

element function *sequence iterator* [Function]
 Returns the sequence element associated to the current position of the iteration.

setf element function *new-value sequence iterator* [Function]
 Destructively modifies the associated sequence by replacing the sequence element associated to the current iteration position with a new value.

index function *sequence iterator* [Function]
 Returns the position of the iteration in the associated sequence.

copy function *sequence iterator* [Function]
 Returns a copy of the iteration state which can be mutated independently of the copied iteration state.

An iterator is created by calling:

sb-sequence:make-sequence-iterator *sequence &key from-end start end* [Generic Function]
 Returns a sequence iterator for *sequence* or, if *start* and/or *end* are supplied, the subsequence bounded by *start* and *end* as nine values:
 1. iterator state 2. limit 3. from-end 4. step function 5. endp function 6. element function 7. setf element function 8. index function 9. copy state function
 If *from-end* is `nil`, the constructed iterator visits the specified elements in the order in which they appear in *sequence*. Otherwise, the elements are visited in the opposite order.

Note that `make-sequence-iterator` calls `make-simple-sequence-iterator` when there is no specialized method for a particular *sequence* subclass. See Section 7.6.2 [Simple Iterator Protocol], page 52.

The following convenience macros simplify traversing sequences using iterators:

sb-sequence:with-sequence-iterator (*&optional iterator limit from-end-p step endp element set-element index copy*) (*sequence &key from-end start end*) *&body body* [Macro]
 Executes *body* with the elements of *vars* bound to the iteration state returned by `make-sequence-iterator` for *sequence* and *args*. Elements of *vars* may be `nil` in which case the corresponding value returned by `make-sequence-iterator` is ignored.

sb-sequence:with-sequence-iterator-functions (*step endp elt setf index copy*) (*sequence &rest args &key from-end start end*) *&body body* [Macro]
 Executes *body* with the names *step*, *endp*, *elt*, *setf*, *index* and *copy* bound to local functions which execute the iteration state query and mutation functions returned by `make-sequence-iterator` for *sequence* and *args*. *step*, *endp*, *elt*, *setf*, *index* and *copy* have dynamic extent.

7.6.2 Simple Iterator Protocol

For cases in which the full flexibility and performance of the general sequence iterator protocol is not required, there is a simplified sequence iterator protocol consisting of a few generic functions which can be specialized for iterator classes:

sb-sequence:iterator-step *sequence iterator from-end* [Generic Function]
 Moves *iterator* one position forward or backward in *sequence* depending on the iteration direction encoded in *from-end*.

sb-sequence:iterator-endp *sequence iterator limit from-end* [Generic Function]
 Returns non-NIL when *iterator* has reached *limit* (which may correspond to the end of *sequence*) with respect to the iteration direction encoded in *from-end*.

sb-sequence:iterator-element *sequence iterator* [Generic Function]
 Returns the element of *sequence* associated to the position of *iterator*.

(setf sb-sequence:iterator-element) [Generic Function]
 Destructively modifies *sequence* by replacing the sequence element associated to position of *iterator* with *new-value*.

sb-sequence:iterator-index *sequence iterator* [Generic Function]
 Returns the position of *iterator* in *sequence*.

sb-sequence:iterator-copy *sequence iterator* [Generic Function]
 Returns a copy of *iterator* which also traverses *sequence* but can be mutated independently of *iterator*.

Iterator objects implementing the above simple iteration protocol are created by calling the following generic function:

sb-sequence:make-simple-sequence-iterator *sequence &key from-end start* [Generic Function]
end
 Returns a sequence iterator for *sequence*, *start*, *end* and *from-end* as three values:
 1. iterator state 2. limit 3. from-end
 The returned iterator can be used with the generic iterator functions *iterator-step*, *iterator-endp*, *iterator-element*, *(setf iterator-element)*, *iterator-index* and *iterator-copy*.

7.7 Support For Unix

7.7.1 Command-line arguments

The UNIX command line can be read from the variable `sb-ext:*posix-argv*`.

7.7.2 Querying the process environment

The UNIX environment can be queried with the `sb-ext:posix-getenv` function.

sb-ext:posix-getenv *name* [Function]
 Return the "value" part of the environment string "name=value" which corresponds to *name*, or nil if there is none.

7.7.3 Running external programs

External programs can be run with `sb-ext:run-program`.²

sb-ext:run-program *program args &key env environment wait search pty input* [Function]
if-input-does-not-exist output if-output-exists error if-error-exists status-hook
external-format directory preserve-fds

`run-program` creates a new process specified by *program*. *args* is a list of strings to be passed literally to the new program. In *posix* environments, this list becomes the array supplied as the second parameter to the `execv()` or `execvp()` system call, each list element becoming one array element. The strings should not contain shell escaping, as there is no shell involvement. Further note that while conventionally the process receives its own pathname in `argv[0]`, that is automatic, and the 0th string should not be present in *args*.

The program arguments and the environment are encoded using the default external format for streams.

² In SBCL versions prior to 1.0.13, `sb-ext:run-program` searched for executables in a manner somewhat incompatible with other languages. As of this version, SBCL uses the system library routine `execvp(3)`, and no longer contains the function, `find-executable-in-search-path`, which implemented the old search. Users who need this function may find it in `run-program.lisp` versions 1.67 and earlier in SBCL's CVS repository here <http://sbcl.cvs.sourceforge.net/sbcl/sbcl/src/code/run-program.lisp?view=log>. However, we caution such users that this search routine finds executables that system library routines do not.

`run-program` will return a `process` structure. See the `cmu` Common Lisp Users Manual for details about the `process` structure.

Notes about Unix environments (as in the `:environment` and `:env` args):

- The `sbc1` implementation of `run-program`, like Perl and many other programs, but unlike the original `cmu cl` implementation, copies the Unix environment by default.
- Running Unix programs from a `setuid` process, or in any other situation where the Unix environment is under the control of someone else, is a mother lode of security problems. If you are contemplating doing this, read about it first. (The Perl community has a lot of good documentation about this and other security issues in script-like programs.)

The &key arguments have the following meanings:

`:environment`

a list of `STRING`s describing the new Unix environment (as in "man environ"). The default is to copy the environment of the current process.

`:env`

an alternative lossy representation of the new Unix environment, for compatibility with `cmu cl`

`:search`

Look for `program` in each of the directories in the child's `$PATH` environment variable. Otherwise an absolute pathname is required.

`:wait`

If non-`NIL` (default), wait until the created process finishes. If `nil`, continue running Lisp until the program finishes.

`:pty` (*not supported on win32*)

Either `t`, `nil`, or a stream. Unless `nil`, the subprocess is established under a `pty`. If `:pty` is a stream, all output to this `pty` is sent to this stream, otherwise the `process-pty` slot is filled in with a stream connected to `pty` that can read output and write input.

`:input`

Either `t`, `nil`, a pathname, a stream, or `:stream`. `t`: the standard input for the current process is inherited. `nil`: `/dev/null` (nul on win32) is used. `pathname`: the specified file is used. `stream`: all the input is read from that stream and sent to the subprocess. `:stream`: the `process-input` slot is filled in with a stream that sends its output to the process. Defaults to `nil`.

`:if-input-does-not-exist` (*when :input is the name of a file*)

can be one of: `:error` to generate an error `:create` to create an empty file `nil` (the default) to return `nil` from `run-program`

`:output`

Either `t`, `nil`, a pathname, a stream, or `:stream`. `t`: the standard output for the current process is inherited. `nil`: `/dev/null` (nul on win32) is used. `pathname`: the specified file is used. `stream`: all the output from the process is written to this stream. `:stream`: the `process-output` slot is filled in with a stream that can be read to get the output. Defaults to `nil`.

`:error`

Same as `:output`, additionally accepts `:output`, making all error output routed to the same place as normal output. Defaults to `:output`.

`:if-output-exists` (*when :output is the name of a file*)

can be one of: `:error` (the default) to generate an error `:supersede` to supersede the file with output from the program `:append` to append output from the program to the file `nil` to return `nil` from `run-program`, without doing anything

`:if-error-exists`

Same as `:if-output-exists`, controlling `:error` output to files. Ignored when `:error` `:output`. Defaults to `:error`.

`:status-hook`

This is a function the system calls whenever the status of the process changes. The function takes the process as an argument.

`:external-format`

The external-format to use for `:input`, `:output`, and `:error` `STREAM`s.

:directory

Specifies the directory in which the program should be run. **nil** (the default) means the directory is unchanged.

:preserve-fds

A sequence of file descriptors which should remain open in the child process.

*Windows specific options:***:escape-arguments** (default **t**)

Controls escaping of the arguments passed to **CreateProcess**.

:window (default **nil**)

When **nil**, the subprocess decides how it will display its window. The following options control how the subprocess window should be displayed: **:hide**, **:show-normal**, **:show-maximized**, **:show-minimized**, **:show-no-activate**, **:show-min-no-active**, **:show-na**. Note: console application subprocesses may or may not display a console window depending on whether the **sbcl** runtime is itself a console or **gui** application. Invoke **cmd /C start** to consistently display a console window or use the **:window :hide** option to consistently hide the console window.

When **sb-ext:run-program** is called with **wait** equal to **NIL**, an instance of class **sb-ext:process** is returned. The following functions are available for use with processes:

sb-ext:process-p *object* [Function]
t if *object* is a **process**, **nil** otherwise.

sb-ext:process-input *instance* [Function]
 The input stream of the process or **nil**.

sb-ext:process-output *instance* [Function]
 The output stream of the process or **nil**.

sb-ext:process-error *instance* [Function]
 The error stream of the process or **nil**.

sb-ext:process-alive-p *process* [Function]
 Return **t** if *process* is still alive, **nil** otherwise.

sb-ext:process-status *process* [Function]
 Return the current status of *process*. The result is one of **:running**, **:stopped**, **:exited**, or **:signaled**.

sb-ext:process-wait *process &optional check-for-stopped* [Function]
 Wait for *process* to quit running for some reason. When *check-for-stopped* is **t**, also returns when *process* is stopped. Returns *process*.

sb-ext:process-exit-code *process* [Function]
 The exit code or the signal of a stopped process.

sb-ext:process-core-dumped *instance* [Function]
t if a core image was dumped by the process.

sb-ext:process-close *process* [Function]
 Close all streams connected to *process* and stop maintaining the status slot.

sb-ext:process-kill *process signal &optional whom* [Function]
 Hand *signal* to *process*. If *whom* is **:pid**, use the **kill** Unix system call. If *whom* is **:process-group**, use the **killpg** Unix system call. If *whom* is **:pty-process-group** deliver the signal to whichever process group is currently in the foreground. Returns **t** if successful, otherwise returns **nil** and error number (two values).

7.8 Unicode Support

SBCL provides support for working with Unicode text and querying the standard Unicode database for information about individual codepoints. Unicode-related functions are located in the `sb-unicode` package.

SBCL also extends ANSI character literal syntax to support Unicode codepoints. You can either specify a character by its Unicode name, with spaces replaced by underscores, if a unique name exists³ or by giving its hexadecimal codepoint preceded by a “U”, an optional “+”, and an arbitrary number of leading zeros. You may also input the character directly into your source code if it can be encoded in your file. If a character had an assigned name in Unicode 1.0 that was distinct from its current name, you may also use that name (with spaces replaced by underscores) to specify the character, unless the name is already associated with a codepoint in the latest Unicode standard (such as “BELL”).

For example, you can specify the codepoint U+00E1 (“Latin Small Letter A With Acute”) as

- `#\LATIN_SMALL_LETTER_A_WITH_ACUTE`
- `#\LATIN_SMALL_LETTER_A_ACUTE`
- `#\á` assuming a Unicode source file
- `#\U00E1`
- `#\UE1`
- `#\U+00E1`

7.8.1 Unicode property access

The following functions can be used to find information about a Unicode codepoint.

`sb-unicode:general-category character` [Function]

Returns the general category of `character` as it appears in `UnicodeData.txt`

`sb-unicode:bidirectional-class character` [Function]

Returns the bidirectional class of `character`

`sb-unicode:combining-class character` [Function]

Returns the canonical combining class (ccc) of `character`

`sb-unicode:decimal-value character` [Function]

Returns the decimal digit value associated with `character` or `nil` if there is no such value.

The only characters in Unicode with a decimal digit value are those that are part of a range of characters that encode the digits 0-9. Because of this, ‘(decimal-digit c) <=> (digit-char-p c 10)’ in `#+sb-unicode` builds

`sb-unicode:digit-value character` [Function]

Returns the Unicode digit value of `character` or `nil` if it doesn’t exist.

Digit values are guaranteed to be integers between 0 and 9 inclusive. All characters with decimal digit values have the same digit value, but there are characters (such as digits of number systems without a 0 value) that have a digit value but no decimal digit value

`sb-unicode:numeric-value character` [Function]

Returns the numeric value of `character` or `nil` if there is no such value. Numeric value is the most general of the Unicode numeric properties. The only constraint on the numeric value is that it be a rational number.

`sb-unicode:mirrored-p character` [Function]

Returns `t` if `character` needs to be mirrored in bidirectional text. Otherwise, returns `nil`.

`sb-unicode:bidirectional-mirroring-glyph character` [Function]

Returns the mirror image of `character` if it exists. Otherwise, returns `nil`.

`sb-unicode:age character` [Function]

Returns the version of Unicode in which `character` was assigned as a pair of values, both integers, representing the major and minor version respectively. If `character` is not assigned in Unicode, returns `nil` for both values.

³ Please note that the codepoint U+1F5CF (PAGE) introduced in Unicode 7.0 is named `UNICODE_PAGE`, since the name “Page” is required to be assigned to form-feed (U+0C) by the ANSI standard.

- `sb-unicode:hangul-syllable-type character` [Function]
Returns the Hangul syllable type of `character`. The syllable type can be one of `:l`, `:v`, `:t`, `:lv`, or `:lvt`. If the character is not a Hangul syllable or Jamo, returns `nil`
- `sb-unicode:east-asian-width character` [Function]
Returns the East Asian Width property of `character` as one of the keywords `:n` (Narrow), `:a` (Ambiguous), `:h` (Halfwidth), `:w` (Wide), `:f` (Fullwidth), or `:na` (Not applicable)
- `sb-unicode:script character` [Function]
Returns the Script property of `character` as a keyword. If `character` does not have a known script, returns `:unknown`
- `sb-unicode:char-block character` [Function]
Returns the Unicode block in which `character` resides as a keyword. If `character` does not have a known block, returns `:no-block`
- `sb-unicode:unicode-1-name character` [Function]
Returns the name assigned to `character` in Unicode 1.0 if it is distinct from the name currently assigned to `character`. Otherwise, returns `nil`. This property has been officially obsoleted by the Unicode standard, and is only included for backwards compatibility.
- `sb-unicode:proplist-p character property` [Function]
Returns `t` if `character` has the specified `property`. `property` is a keyword representing one of the properties from PropList.txt, with underscores replaced by dashes.
- `sb-unicode:uppercase-p character` [Function]
Returns `t` if `character` has the Unicode property Uppercase and `nil` otherwise
- `sb-unicode:lowercase-p character` [Function]
Returns `t` if `character` has the Unicode property Lowercase and `nil` otherwise
- `sb-unicode:cased-p character` [Function]
Returns `t` if `character` has a (Unicode) case, and `nil` otherwise
- `sb-unicode:case-ignorable-p character` [Function]
Returns `t` if `character` is Case Ignorable as defined in Unicode 6.3, Chapter 3
- `sb-unicode:alphabetic-p character` [Function]
Returns `t` if `character` is Alphabetic according to the Unicode standard and `nil` otherwise
- `sb-unicode:ideographic-p character` [Function]
Returns `t` if `character` has the Unicode property Ideographic, which loosely corresponds to the set of "Chinese characters"
- `sb-unicode:math-p character` [Function]
Returns `t` if `character` is a mathematical symbol according to Unicode and `nil` otherwise
- `sb-unicode:whitespace-p character` [Function]
Returns `t` if `character` is whitespace according to Unicode and `nil` otherwise
- `sb-unicode:soft-dotted-p character` [Function]
Returns `t` if `character` has a soft dot (such as the dots on `i` and `j`) which disappears when accents are placed on top of it. and `nil` otherwise
- `sb-unicode:hex-digit-p character &key ascii` [Function]
Returns `t` if `character` is a hexadecimal digit and `nil` otherwise. If `:ascii` is non-NIL, fullwidth equivalents of the Latin letters `A` through `f` are excluded.
- `sb-unicode:default-ignorable-p character` [Function]
Returns `t` if `character` is a Default_Ignorable_Code_Point
- `sb-unicode:grapheme-break-class char` [Function]
Returns the grapheme breaking class of `character`, as specified in uax #29.
- `sb-unicode:word-break-class char` [Function]
Returns the word breaking class of `character`, as specified in uax #29.

sb-unicode:sentence-break-class *char* [Function]
Returns the sentence breaking class of **character**, as specified in **uax #29**.

sb-unicode:line-break-class *character &key resolve* [Function]
Returns the line breaking class of **character**, as specified in **uax #14**. If **:resolve** is **nil**, returns the character class found in the property file. If **:resolve** is non-NIL, certain line-breaking classes will be mapped to other classes as specified in the applicable standards. Additionally, if **:resolve** is **:east-asian**, Ambiguous (class **:ai**) characters will be mapped to the Ideographic (**:id**) class instead of Alphabetic (**:al**).

7.8.2 String operations

SBCL can normalize strings using:

sb-unicode:normalize-string *string &optional form filter* [Function]
Normalize **string** to the Unicode normalization form form. Acceptable values for form are **:nfd**, **:nfc**, **:nfkd**, and **:nfkc**. If **filter** is a function it is called on each decomposed character and only characters for which it returns **t** are collected.

sb-unicode:normalized-p *string &optional form* [Function]
Tests if **string** is normalized to **form**

SBCL implements the full range of Unicode case operations with the functions

sb-unicode:uppercase *string &key locale* [Function]
Returns the full uppercase of **string** according to the Unicode standard. The result is not guaranteed to have the same length as the input. If **:locale** is **nil**, no language-specific case transformations are applied. If **:locale** is a keyword representing a two-letter **iso** country code, the case transforms of that locale are used. If **:locale** is **t**, the user's current locale is used (Unix and Win32 only).

sb-unicode:lowercase *string &key locale* [Function]
Returns the full lowercase of **string** according to the Unicode standard. The result is not guaranteed to have the same length as the input. **:locale** has the same semantics as the **:locale** argument to **uppercase**.

sb-unicode:titlecase *string &key locale* [Function]
Returns the titlecase of **string**. The resulting string can be longer than the input. **:locale** has the same semantics as the **:locale** argument to **uppercase**.

sb-unicode:casefold *string* [Function]
Returns the full casefolding of **string** according to the Unicode standard. Casefolding removes case information in a way that allows the results to be used for case-insensitive comparisons. The result is not guaranteed to have the same length as the input.

It also extends standard Common Lisp case functions such as **string-upcase** and **string-downcase** to support a subset of Unicode's casing behavior. Specifically, a character is **both-case-p** if its case mapping in Unicode is one-to-one and invertable.

The **sb-unicode** package also provides functions for collating/sorting strings according to the Unicode Collation Algorithm.

sb-unicode:unicode< *string1 string2 &key start1 end1 start2 end2* [Function]
Determines whether **STRING1** sorts before **STRING2** using the Unicode Collation Algorithm. The function uses an untailored Default Unicode Collation Element Table to produce the sort keys. The function uses the Shifted method for dealing with variable-weight characters, as described in **uts #10**

sb-unicode:unicode= *string1 string2 &key start1 end1 start2 end2 strict* [Function]
Determines whether **STRING1** and **STRING2** are canonically equivalent according to Unicode. The **start** and **end** arguments behave like the arguments to **STRING=**. If **:strict** is **nil**, **UNICODE=** tests compatibility equivalence instead.

sb-unicode:unicode-equal *string1 string2 &key start1 end1 start2 end2 strict* [Function]
Determines whether **STRING1** and **STRING2** are canonically equivalent after casefoldin8 (that is, ignoring case differences) according to Unicode. The **start** and **end** arguments behave like the arguments to **STRING=**. If **:strict** is **nil**, **UNICODE=** tests compatibility equivalence instead.

sb-unicode:unicode<= *string1 string2 &key start1 end1 start2 end2* [Function]
 Tests if STRING1 and STRING2 are either UNICODE< or UNICODE=

sb-unicode:unicode> *string1 string2 &key start1 end1 start2 end2* [Function]
 Tests if STRING2 is UNICODE< STRING1.

sb-unicode:unicode>= *string1 string2 &key start1 end1 start2 end2* [Function]
 Tests if STRING1 and STRING2 are either UNICODE= or UNICODE>

The following functions are provided for detecting visually confusable strings:

sb-unicode:confusable-p *string1 string2 &key start1 end1 start2 end2* [Function]
 Determines whether STRING1 and STRING2 could be visually confusable according to the `idna-confusableSummary.txt` table

7.8.3 Breaking strings

The **sb-unicode** package includes several functions for breaking a Unicode string into useful parts.

sb-unicode:graphemes *string* [Function]
 Breaks **string** into graphemes according to the default grapheme breaking rules specified in **uax #29**, returning a list of strings.

sb-unicode:words *string* [Function]
 Breaks **string** into words according to the default word breaking rules specified in **uax #29**. Returns a list of strings

sb-unicode:sentences *string* [Function]
 Breaks **string** into sentences according to the default sentence breaking rules specified in **uax #29**

sb-unicode:lines *string &key margin* [Function]
 Breaks **string** into lines that are no wider than **:margin** according to the line breaking rules outlined in **uax #14**. Combining marks will always be kept together with their base characters, and spaces (but not other types of whitespace) will be removed from the end of lines. If **:margin** is unspecified, it defaults to 80 characters

7.9 Customization Hooks for Users

The toplevel repl prompt may be customized, and the function that reads user input may be replaced completely.

The behaviour of **require** when called with only one argument is implementation-defined. In SBCL, **require** behaves in the following way:

cl:require *module-name &optional pathnames* [Function]
 Loads a module, unless it already has been loaded. **pathnames**, if supplied, is a designator for a list of pathnames to be loaded if the module needs to be. If **pathnames** is not supplied, functions from the list ***module-provider-functions*** are called in order with **module-name** as an argument, until one of them returns non-NIL. User code is responsible for calling **provide** to indicate a successful load of the module.

sb-ext:*module-provider-functions* [Variable]
 See function documentation for **require**.

Although SBCL does not provide a resident editor, the **ed** function can be customized to hook into user-provided editing mechanisms as follows:

cl:ed *&optional x* [Function]
 Starts the editor (on a file or a function if named). Functions from the list ***ed-functions*** are called in order with **x** as an argument until one of them returns non-NIL; these functions are responsible for signalling a **file-error** to indicate failure to perform an operation on the file system.

sb-ext:*ed-functions* [Variable]
 See function documentation for **ed**.

Conditions of type **warning** and **style-warning** are sometimes signaled at runtime, especially during execution of Common Lisp defining forms such as **defun**, **defmethod**, etc. To muffle these warnings at runtime, SBCL provides a variable **sb-ext:*muffled-warnings***:

sb-ext:*muffled-warnings* [Variable]

A type that ought to specify a subtype of **warning**. Whenever a warning is signaled, if the warning is of this type and is not handled by any other handler, it will be muffled.

7.10 Tools To Help Developers

SBCL provides a profiler and other extensions to the ANSI **trace** facility. For more information, see [Macro **common-lisp trace**], page 34.

The debugger supports a number of options. Its documentation is accessed by typing **help** at the debugger prompt. See Chapter 5 [Debugger], page 26.

Documentation for **inspect** is accessed by typing **help** at the **inspect** prompt.

7.11 Resolution of Name Conflicts

The ANSI standard (section 11.1.1.2.5) requires that name conflicts in packages be resolvable in favour of any of the conflicting symbols. In the interactive debugger, this is achieved by prompting for the symbol in whose favour the conflict should be resolved; for programmatic use, the **sb-ext:resolve-conflict** restart should be invoked with one argument, which should be a member of the list returned by the condition accessor **sb-ext:name-conflict-symbols**.

7.12 Hash Table Extensions

Hash table extensions supported by SBCL are all controlled by keyword arguments to **make-hash-table**.

cl:make-hash-table *&key test size rehash-size rehash-threshold* [Function]
hash-function weakness synchronized

Create and return a new hash table. The keywords are as follows:

- :test** Determines how keys are compared. Must a designator for one of the standard hash table tests, or a hash table test defined using **sb-ext:define-hash-table-test**. Additionally, when an explicit **hash-function** is provided (see below), any two argument equivalence predicate can be used as the **test**.
- :size** A hint as to how many elements will be put in this hash table.
- :rehash-size** Indicates how to expand the table when it fills up. If an integer, add space for that many elements. If a floating point number (which must be greater than 1.0), multiply the size by that amount.
- :rehash-threshold** Indicates how dense the table can become before forcing a rehash. Can be any positive number ≤ 1 , with density approaching zero as the threshold approaches 0. Density 1 means an average of one entry per bucket.
- :hash-function** If unsupplied, a hash function based on the **test** argument is used, which then must be one of the standardized hash table test functions, or one for which a default hash function has been defined using **sb-ext:define-hash-table-test**. If **hash-function** is specified, the **test** argument can be any two argument predicate consistent with it. The **hash-function** is expected to return a non-negative fixnum hash code. If **test** is neither standard nor defined by **define-hash-table-test**, then the **hash-function** must be specified.
- :weakness** When **:weakness** is not **nil**, garbage collection may remove entries from the hash table. The value of **:weakness** specifies how the presence of a key or value in the hash table preserves their entries from garbage collection.
 Valid values are:
:key means that the key of an entry must be live to guarantee that the entry is preserved.

:value means that the value of an entry must be live to guarantee that the entry is preserved.

:key-and-value means that both the key and the value must be live to guarantee that the entry is preserved.

:key-or-value means that either the key or the value must be live to guarantee that the entry is preserved.

nil (the default) means that entries are always preserved.

:synchronized

If **nil** (the default), the hash-table may have multiple concurrent readers, but results are undefined if a thread writes to the hash-table concurrently with another reader or writer. If **t**, all concurrent accesses are safe, but note that **clhs 3.6** (Traversal Rules and Side Effects) remains in force. See also: **sb-ext:with-locked-hash-table**. This keyword argument is experimental, and may change incompatibly or be removed in the future.

sb-ext:define-hash-table-test *name hash-function*

[Macro]

Defines **name** as a new kind of hash table test for use with the **:test** argument to **make-hash-table**, and associates a default **hash-function** with it.

name must be a symbol naming a global two argument equivalence predicate. Afterwards both **'name** and **#'name** can be used with **:test** argument. In both cases **hash-table-test** will return the symbol **name**.

hash-function must be a symbol naming a global hash function consistent with the predicate, or be a **lambda** form implementing one in the current lexical environment. The hash function must compute the same hash code for any two objects for which **name** returns true, and subsequent calls with already hashed objects must always return the same hash code.

Note: The **:hash-function** keyword argument to **make-hash-table** can be used to override the specified default hash-function.

Attempting to define **name** in a locked package as hash-table test causes a package lock violation.

Examples:

```
;;; 1.
```

```
;; We want to use objects of type FOO as keys (by their
;; names.) EQUALP would work, but would make the names
;; case-insensitive -- which we don't want.
(defstruct foo (name nil :type (or null string)))

;; Define an equivalence test function and a hash function.
(defun foo-name= (f1 f2) (equal (foo-name f1) (foo-name f2)))
(defun sxhash-foo-name (f) (sxhash (foo-name f)))

(define-hash-table-test foo-name= sxhash-foo-name)

;; #'foo-name would work too.
(defun make-foo-table () (make-hash-table :test 'foo-name=))
```

```
;;; 2.
```

```
(defun == (x y) (= x y))

(define-hash-table-test ==
  (lambda (x)
    ;; Hash codes must be consistent with test, so
    ;; not (SXHASH X), since
    ;;   (= 1 1.0) => T
    ;;   (= (SXHASH 1) (SXHASH 1.0)) => NIL
    ;; Note: this doesn't deal with complex numbers or
    ;; bignums too large to represent as double floats.
    (sxhash (coerce x 'double-float)))))
```

```
;; #'== would work too
(defun make-number-table () (make-hash-table :test '==))
```

sb-ext:with-locked-hash-table (*hash-table*) &*body* *body* [Macro]
Limits concurrent accesses to **hash-table** for the duration of *body*. If **hash-table** is synchronized, *body* will execute with exclusive ownership of the table. If **hash-table** is not synchronized, *body* will execute with other **with-locked-hash-table** bodies excluded -- exclusion of hash-table accesses not surrounded by **with-locked-hash-table** is unspecified.

sb-ext:hash-table-synchronized-p *ht* [Function]
Returns **t** if **hash-table** is synchronized.

sb-ext:hash-table-weakness *ht* [Function]
Return the weakness of **hash-table** which is one of **nil**, **:key**, **:value**, **:key-and-value**, **:key-or-value**.

7.13 Random Number Generation

The initial value of ***random-state*** is the same each time SBCL is started. This makes it possible for user code to obtain repeatable pseudo random numbers using only standard-provided functionality. See **seed-random-state** below for an SBCL extension that allows to seed the random number generator from given data for an additional possibility to achieve this. Non-repeatable random numbers can always be obtained using (**make-random-state** **t**).

The sequence of numbers produced by repeated calls to **random** starting with the same random state and using the same sequence of **limit** arguments is guaranteed to be reproducible only in the same version of SBCL on the same platform, using the same code under the same evaluator mode and compiler optimization qualities. Just two examples of differences that may occur otherwise: calls to **random** can be compiled differently depending on how much is known about the **limit** argument at compile time, yielding different results even if called with the same argument at run time, and the results can differ depending on the machine's word size, for example for limits that are fixnums under 64-bit word size but bignums under 32-bit word size.

sb-ext:seed-random-state &*optional state* [Function]
Make a random state object. The optional **state** argument specifies a seed for deterministic pseudo-random number generation.

As per the Common Lisp standard for **make-random-state**,

- If **state** is **nil** or not supplied, return a copy of the default ***random-state***.
- If **state** is a random state, return a copy of it.
- If **state** is **t**, return a randomly initialized state (using operating-system provided randomness where available, otherwise a poor substitute based on internal time and pid).

As a supported **sbcl** extension, we also support receiving as a seed an object of the following types:

- (**simple-array** (**unsigned-byte** 8) (*))
- **unsigned-byte**

While we support arguments of any size and will mix the provided bits into the random state, it is probably overkill to provide more than 256 bits worth of actual information.

This particular **sbcl** version also accepts an argument of the following type: (**simple-array** (**unsigned-byte** 32) (*))

This particular **sbcl** version uses the popular MT19937 **prng** algorithm, and its internal state only effectively contains about 19937 bits of information. <http://www.math.sci.hiroshima-u.ac.jp/~m-mat/MT/emt.html>

Some notes on random floats: The standard doesn't prescribe a specific method of generating random floats. The following paragraph describes SBCL's current implementation and should be taken as purely informational, that is, user code should not depend on any of its specific properties. The method used has been chosen because it is common, conceptually simple and fast.

To generate random floats, SBCL evaluates code that has an equivalent effect as

```
(* limit
```

```
(float (/ (random (expt 2 23)) (expt 2 23)) 1.0f0))
```

(for single-floats) and correspondingly (with 52 and 1.0d0 instead of 23 and 1.0f0) for double-floats. Note especially that this means that zero is a possible return value occurring with probability (expt 2 -23) respectively (expt 2 -52). Also note that there exist twice as many equidistant floats between 0 and 1 as are generated. For example, the largest number that (random 1.0f0) ever returns is (float (/ (1- (expt 2 23)) (expt 2 23)) 1.0f0) while (float (/ (1- (expt 2 24)) (expt 2 24)) 1.0f0) is the largest single-float less than 1. This is a side effect of the fact that the implementation uses the fastest possible conversion from bits to floats.

SBCL currently uses the Mersenne Twister as its random number generator, specifically the 32-bit version under both 32- and 64-bit word size. The seeding algorithm has been improved several times by the authors of the Mersenne Twister; SBCL uses the third version (from 2002) which is still the most recent as of June 2012. The implementation has been tested to provide output identical to the recommended C implementation.

While the Mersenne Twister generates random numbers of much better statistical quality than other widely used generators, it uses only linear operations modulo 2 and thus fails some statistical tests⁴. For example, the distribution of ranks of (sufficiently large) random binary matrices is much distorted compared to the theoretically expected one when the matrices are generated by the Mersenne Twister. Thus, applications that are sensitive to this aspect should use a different type of generator.

7.14 Timeouts and Deadlines

SBCL supports three different ways of restricting the execution time available to individual operations or parts of computations:

Timeout Parameters

Some operations such as thread synchronization primitives accept a `:timeout` parameter. See Section 7.14.1 [Timeout Parameters], page 63.

Synchronous Timeouts (Deadlines)

Certain operations that may suspend execution for extended periods of time such as `cl:sleep`, thread synchronization primitives, IO and waiting for external processes respect deadlines established for a part of a computation. See Section 7.14.2 [Synchronous Timeouts (Deadlines)], page 64.

Asynchronous Timeouts

Asynchronous timeouts can interrupt most computations at (almost) any point. Thus, this kind of timeouts is the most versatile but it is also somewhat unsafe. See Section 7.14.3 [Asynchronous Timeouts], page 64.

7.14.1 Timeout Parameters

Certain operations accept `:timeout` keyword arguments. These only affect the specific operation and must be specified at each call site by passing a `:timeout` keyword argument and a corresponding timeout value to the respective operation. Expiration of the timeout before the operation completes results in either a normal return with a return value indicating the timeout or in the signaling of a specialized condition such as `sb-thread:join-thread-error`.

Example:

```
(defun join-thread-within-5-seconds (thread)
  (multiple-value-bind (value result)
    (sb-thread:join-thread thread :default nil :timeout 5)
    (when (eq result :timeout)
      (error "Could not join ~A within 5 seconds" thread))
    value))
```

The above code attempts to join the specified thread for up to five seconds, returning its value in case of success. If the thread is still running after the five seconds have elapsed, `sb-thread:join-thread` indicates the timeout in its second return value. If a `:default` value was not provided, `sb-thread:join-thread` would signal a `sb-thread:join-thread-error` instead.

To wait for an arbitrary condition, optionally with a timeout, the `sb-ext:wait-for` macro can be used:

⁴ See chapter 7 "Testing widely used RNGs" in *TestU01: A C Library for Empirical Testing of Random Number Generators* by Pierre L'Ecuyer and Richard Simard, ACM Transactions on Mathematical Software, Vol. 33, article 22, 2007.

sb-ext:wait-for *test-form &key timeout*

[Macro]

Wait until **test-form** evaluates to true, then return its primary value. If **timeout** is provided, waits at most approximately **timeout** seconds before returning **nil**.

If **with-deadline** has been used to provide a global deadline, signals a **deadline-timeout** if **test-form** doesn't evaluate to true before the deadline.

Experimental: subject to change without prior notice.

7.14.2 Synchronous Timeouts (Deadlines)

Deadlines, in contrast to timeout parameters, are established for a dynamic scope using the **sb-sys:with-deadline** macro and indirectly affect operations within that scope. In case of nested uses, the effective deadline is the one that expires first unless an inner use explicitly overrides outer deadlines.

sb-sys:with-deadline (*&key seconds override*) *&body body*

[Macro]

Arranges for a **timeout** condition to be signalled if an operation respecting deadlines occurs either after the deadline has passed, or would take longer than the time left to complete.

Currently only **sleep**, blocking **io** operations, **get-mutex**, and **condition-wait** respect deadlines, but this includes their implicit uses inside **sbcl** itself.

Unless **override** is true, existing deadlines can only be restricted, not extended. Deadlines are per thread: children are unaffected by their parent's deadlines.

Experimental.

Expiration of deadlines set up this way only has an effect when it happens before or during the execution of a deadline-aware operation (see Section 7.14.4 [Operations Supporting Timeouts and Deadlines], page 65). In this case, a **sb-sys:deadline-timeout** is signaled. A handler for this condition type may use the **sb-sys:defer-deadline** or **sb-sys:cancel-deadline** restarts to defer or cancel the deadline respectively and resume execution of the interrupted operation.

sb-sys:deadline-timeout

[Condition]

Class precedence list: **deadline-timeout**, **timeout**, **serious-condition**, **condition**, **t**

Signaled when an operation in the context of a deadline takes longer than permitted by the deadline.

When a thread is executing the debugger, signaling of **sb-sys:deadline-timeout** conditions for that thread is deferred until it exits the debugger.

Example:

```
(defun read-input ()
  (list (read-line) (read-line)))

(defun do-it ()
  (sb-sys:with-deadline (:seconds 5))
    (read-input)
    (sleep 2)
    (sb-ext:run-program "my-program"))
```

The above code establishes a deadline of five seconds within which the body of the **do-it** function should execute. All calls of deadline-aware functions in the dynamic scope, in this case two **read-line** calls, a **sleep** call and a **sb-ext:run-program** call, are affected by the deadline. If, for example, the first **read-line** call completes in one second and the second **read-line** call completes in three seconds, a **sb-sys:deadline-timeout** condition will be signaled after the **sleep** call has been executing for one second.

7.14.3 Asynchronous Timeouts

Asynchronous timeouts are established for a dynamic scope using the **sb-sys:with-timeout** macro:

sb-ext:with-timeout *expires &body body*

[Macro]

Execute the body, asynchronously interrupting it and signalling a **timeout** condition after at least **expires** seconds have passed.

Note that it is never safe to unwind from an asynchronous condition. Consider:

```
(defun call-with-foo (function)
  (let (foo)
    (unwind-protect
```



```
(progn
  (setf foo (get-foo))
  (funcall function foo))
(when foo
  (release-foo foo))))
```

If `timeout` occurs after `get-foo` has executed, but before the assignment, then `release-foo` will be missed. While individual sites like this can be made proof against asynchronous unwinds, this doesn't solve the fundamental issue, as all the frames potentially unwound through need to be proofed, which includes both system and application code -- and in essence proofing everything will make the system uninterruptible.

Expiration of the timeout will cause the operation being executed at that moment to be interrupted by an asynchronously signaled `sb-ext:timeout` condition, (almost) irregardless of the operation and its context.

`sb-ext:timeout` [Condition]

Class precedence list: `timeout`, `serious-condition`, `condition`, `t`

Signaled when an operation does not complete within an allotted time budget.

7.14.4 Operations Supporting Timeouts and Deadlines

Operation	Timeout Parameter	Affected by Deadlines
<code>cl:sleep</code>	-	since SBCL 1.4.3
<code>cl:read-line</code> , etc.	no	yes
<code>[wait-for]</code> , page 63	yes	yes
<code>[process-wait]</code> , page 55	no	yes
<code>[grab-mutex]</code> , page 102	yes	yes
<code>[condition-wait]</code> , page 104	yes	yes
<code>[wait-on-semaphore]</code> , page 103	yes	yes
<code>[join-thread]</code> , page 97	yes	yes
<code>[receive-message]</code> , page 123	yes	yes?
<code>[wait-on-gate]</code> , page 124	yes	yes?
<code>[frlock-write]</code> , page 125	yes	yes?
<code>[grab-frlock-write-lock]</code> , page 126	yes	yes?

7.15 Miscellaneous Extensions

`sb-ext:array-storage-vector array` [Function]

Returns the underlying storage vector of `array`, which must be a non-displaced array.

In `sbcl`, if `array` is of type `(simple-array * (*))`, it is its own storage vector. Multidimensional arrays, arrays with fill pointers, and adjustable arrays have an underlying storage vector with the same `array-element-type` as `array`, which this function returns.

Important note: the underlying vector is an implementation detail. Even though this function exposes it, changes in the implementation may cause this function to be removed without further warning.

`sb-ext:delete-directory pathspec &key recursive` [Function]

Deletes the directory designated by `pathspec` (a pathname designator). Returns the truename of the directory deleted.

If `recursive` is false (the default), signals an error unless the directory is empty. If `recursive` is true, first deletes all files and subdirectories. If `recursive` is true and the directory contains symbolic links, the links are deleted, not the files and directories they point to.

Signals an error if `pathspec` designates a file or a symbolic link instead of a directory, or if the directory could not be deleted for any reason.

Both

```
(DELETE-DIRECTORY "/tmp/foo")
(DELETE-DIRECTORY "/tmp/foo/")
```

delete the "foo" subdirectory of "/tmp", or signal an error if it does not exist or if it is a file or a symbolic link.

`sb-ext:get-time-of-day`

[Function]

Return the number of seconds and microseconds since the beginning of the `unix` epoch (January 1st 1970.)

`sb-ext:assert-version->= &rest subversions`

[Function]

Asserts that the current `sbcl` is of version equal to or greater than the version specified in the arguments. A continuable error is signaled otherwise.

The arguments specify a sequence of subversion numbers in big endian order. They are compared lexicographically with the runtime version, and versions are treated as though trailed by an unbounded number of 0s.

For example, `(assert-version->= 1 1 4)` asserts that the current `sbcl` is version 1.1.4[.0.0...] or greater, and `(assert-version->= 1)` that it is version 1[.0.0...] or greater.

7.16 Stale Extensions

SBCL has inherited from CMUCL various hooks to allow the user to tweak and monitor the garbage collection process. These are somewhat stale code, and their interface might need to be cleaned up. If you have urgent need of them, look at the code in `src/code/gc.lisp` and bring it up on the developers' mailing list.

SBCL has various hooks inherited from CMUCL, like `sb-ext:float-denormalized-p`, to allow a program to take advantage of IEEE floating point arithmetic properties which aren't conveniently or efficiently expressible using the ANSI standard. These look good, and their interface looks good, but IEEE support is slightly broken due to a stupid decision to remove some support for infinities (because it wasn't in the ANSI spec and it didn't occur to me that it was in the IEEE spec). If you need this stuff, take a look at the code and bring it up on the developers' mailing list.

7.17 Efficiency Hacks

The `sb-ext:purify` function causes SBCL first to collect all garbage, then to mark all uncollected objects as permanent, never again attempting to collect them as garbage. This can cause a large increase in efficiency when using a primitive garbage collector, or a more moderate increase in efficiency when using a more sophisticated garbage collector which is well suited to the program's memory usage pattern. It also allows permanent code to be frozen at fixed addresses, a precondition for using copy-on-write to share code between multiple Lisp processes. This is less important with modern generational garbage collectors, but not all SBCL platforms use such a garbage collector.

The `sb-ext:truly-the` special form declares the type of the result of the operations, producing its argument; the declaration is not checked. In short: don't use it.

`sb-ext:truly-the value-type form`

[Special Operator]

Specifies that the values returned by `form` conform to the `value-type`, and causes the compiler to trust this information unconditionally.

Consequences are undefined if any result is not of the declared type -- typical symptoms including memory corruptions. Use with great care.

The `sb-ext:freeze-type` declaration declares that a type will never change, which can make type testing (`typep`, etc.) more efficient for structure types.

8 External Formats

External formats determine the coding of characters from/to sequences of octets when exchanging data with the outside world. Examples of such exchanges are:

1. Character streams associated with files, sockets and process input/output (See Section 11.1 [Stream External Formats], page 85, and Section 7.7.3 [Running external programs], page 53)
2. Names of files
3. Foreign strings (See Section 9.2.2 [Foreign Types and Lisp Types], page 72)
4. Posix interface (See Section 17.6 [sb-posix], page 133)
5. Hostname- and protocol-related functions of the BSD-socket interface (See Chapter 15 [Networking], page 110)

Technically, external formats in SBCL are named objects describing coding of characters as well as policies in case de- or encoding is not possible. Each external format has a canonical name and zero or more aliases. User code mostly interacts with external formats by supplying external format designators to functions that use external formats internally.

8.1 The Default External Format

Most functions interacting with external formats use a default external format if none is explicitly specified. In some cases, the default external format is used unconditionally.

The default external format is UTF-8. It can be changed via

`sb-ext:*default-external-format*` and `sb-ext:*default-c-string-external-format*`

8.2 External Format Designators

In situations where an external format designator is required, such as the `:external-format` argument in calls to `open` or `with-open-file`, users may supply the name of an encoding to denote the external format which is applying that encoding to Lisp characters.

In addition to the basic encoding for an external format, options controlling various special cases may be passed, by using a list (whose first element must be an encoding name and whose rest is a plist) as an external file format designator.

More specifically, external format designators can take the following forms:

:default Designates the current default external format (See Section 8.1 [The Default External Format], page 67).

keyword Designates the supported external format that has *keyword* as one of its names. (See Section 8.5 [Supported External Formats], page 68).

(keyword :replacement replacement)

Designates an external format that is like the one designated by *keyword* but does not signal an error in case a character or octet sequence cannot be en- or decoded. Instead, it inserts *replacement* at the position in question. *replacement* has to be a string designator, that is a character or string.

For example:

```
(with-open-file (stream pathname :external-format '(:utf-8 :replacement #\?))
  (read-line stream))
```

will read the first line of *pathname*, replacing any octet sequence that is not valid in the UTF-8 external format with a question mark character.

8.3 Character Coding Conditions

De- or encoding characters using a given external format is not always possible:

- Decoding an octet vector using a given external format can fail if it contains an octet or sequence of octets that does not have an interpretation as a character according to the external format.
- Conversely, a string may contain characters that a given external format cannot encode. For example, the ASCII external format cannot encode the character `#\ö`.

Unless the external format governing the coding uses the `:replacement` keyword, SBCL will signal (continuable) errors under the above circumstances. The types of the condition signaled are not currently exported or documented but will be in future SBCL versions.

8.4 Converting between Strings and Octet Vectors

To encode Lisp strings as octet vectors and decode octet vectors as Lisp strings, the following SBCL-specific functions can be used:

`sb-ext:string-to-octets` *string &key external-format start end null-terminate* [Function]

Return an octet vector that is *string* encoded according to *external-format*.

If *external-format* is given, it must designate an external format.

If given, *start* and *end* must be bounding index designators and designate a subsequence of *string* that should be encoded.

If *null-terminate* is true, the returned octet vector ends with an additional 0 element that does not correspond to any part of *string*.

If some of the characters of *string* (or the subsequence bounded by *start* and *end*) cannot be encoded by *external-format* an error of a subtype of `sb-int:character-encoding-error` is signaled.

Note that for some values of *external-format* and *null-terminate* the length of the returned vector may be different from the length of *string* (or the subsequence bounded by *start* and *end*).

`sb-ext:octets-to-string` *vector &key external-format start end* [Function]

Return a string obtained by decoding *vector* according to *external-format*.

If *external-format* is given, it must designate an external format.

If given, *start* and *end* must be bounding index designators and designate a subsequence of *vector* that should be decoded.

If some of the octets of *vector* (or the subsequence bounded by *start* and *end*) cannot be decoded by *external-format* an error of a subtype of `sb-int:character-decoding-error` is signaled.

Note that for some values of *external-format* the length of the returned string may be different from the length of *vector* (or the subsequence bounded by *start* and *end*).

8.5 Supported External Formats

The following table lists the external formats supported by SBCL in the form of the respective canonical name followed by the list of aliases:

<code>:ASCII</code>	<code>:US-ASCII</code> , <code>:ANSI_X3.4-1968</code> , <code>:ISO-646</code> , <code>:ISO-646-US</code> , <code>: 646 </code>
<code>:CP1250</code>	<code>: cp1250 </code> , <code>:WINDOWS-1250</code> , <code>: windows-1250 </code>
<code>:CP1251</code>	<code>: cp1251 </code> , <code>:WINDOWS-1251</code> , <code>: windows-1251 </code>
<code>:CP1252</code>	<code>: cp1252 </code> , <code>:WINDOWS-1252</code> , <code>: windows-1252 </code>
<code>:CP1253</code>	<code>: cp1253 </code> , <code>:WINDOWS-1253</code> , <code>: windows-1253 </code>
<code>:CP1254</code>	<code>: cp1254 </code>
<code>:CP1255</code>	<code>: cp1255 </code> , <code>:WINDOWS-1255</code> , <code>: windows-1255 </code>
<code>:CP1256</code>	<code>: cp1256 </code> , <code>:WINDOWS-1256</code> , <code>: windows-1256 </code>
<code>:CP1257</code>	<code>: cp1257 </code> , <code>:WINDOWS-1257</code> , <code>: windows-1257 </code>
<code>:CP1258</code>	<code>: cp1258 </code> , <code>:WINDOWS-1258</code> , <code>: windows-1258 </code>
<code>:CP437</code>	<code>: cp437 </code>
<code>:CP850</code>	<code>: cp850 </code>
<code>:CP852</code>	<code>: cp852 </code>
<code>:CP855</code>	<code>: cp855 </code>
<code>:CP857</code>	<code>: cp857 </code>

```

:CP860      :|cp860|
:CP861      :|cp861|
:CP862      :|cp862|
:CP863      :|cp863|
:CP864      :|cp864|
:CP865      :|cp865|
:CP866      :|cp866|
:CP869      :|cp869|
:CP874      :|cp874|
:EBCDIC-US
      :CP037, :|cp037|, :IBM-037, :IBM037
:EUC-JP     :EUCJP, :|eucJP|
:GBK        :CP936
:ISO-8859-10
      :|iso-8859-10|, :LATIN-6, :|latin-6|
:ISO-8859-11
      :|iso-8859-11|
:ISO-8859-13
      :|iso-8859-13|, :LATIN-7, :|latin-7|
:ISO-8859-14
      :|iso-8859-14|, :LATIN-8, :|latin-8|
:ISO-8859-2
      :|iso-8859-2|, :LATIN-2, :|latin-2|
:ISO-8859-3
      :|iso-8859-3|, :LATIN-3, :|latin-3|
:ISO-8859-4
      :|iso-8859-4|, :LATIN-4, :|latin-4|
:ISO-8859-5
      :|iso-8859-5|
:ISO-8859-6
      :|iso-8859-6|
:ISO-8859-7
      :|iso-8859-7|
:ISO-8859-8
      :|iso-8859-8|
:ISO-8859-9
      :|iso-8859-9|, :LATIN-5, :|latin-5|
:KOI8-R     :|koi8-r|
:KOI8-U     :|koi8-u|
:LATIN-1    :LATIN1, :ISO-8859-1, :ISO8859-1
:LATIN-9    :LATIN9, :ISO-8859-15, :ISO8859-15
:MAC-ROMAN
      :|mac-roman|, :|MacRoman|, :MAC, :|mac|, :MACINTOSH, :|macintosh|

```

```
:SHIFT_JIS
    :SJIS, :|Shift_JIS|, :CP932
:UCS-2BE    :UCS2BE
:UCS-2LE    :UCS2LE
:UCS-4BE    :UCS4BE
:UCS-4LE    :UCS4LE
:UTF-16BE   :UTF16BE
:UTF-16LE   :UTF16LE
:UTF-32BE   :UTF32BE
:UTF-32LE   :UTF32LE
:X-MAC-CYRILLIC
    :|x-mac-cyrillic|
```

9 Foreign Function Interface

This chapter describes SBCL’s interface to C programs and libraries (and, since C interfaces are a sort of *lingua franca* of the Unix world, to other programs and libraries in general.)

Note: In the modern Lisp world, the usual term for this functionality is Foreign Function Interface, or FFI, where despite the mention of “function” in this term, FFI also refers to direct manipulation of C data structures as well as functions. The traditional CMUCL terminology is Alien Interface, and while that older terminology is no longer used much in the system documentation, it still reflected in names in the implementation, notably in the name of the SB-ALIEN package.

9.1 Introduction to the Foreign Function Interface

Because of Lisp’s emphasis on dynamic memory allocation and garbage collection, Lisp implementations use non-C-like memory representations for objects. This representation mismatch creates friction when a Lisp program must share objects with programs which expect C data. There are three common approaches to establishing communication:

- The burden can be placed on the foreign program (and programmer) by requiring the knowledge and use of the representations used internally by the Lisp implementation. This can require a considerable amount of “glue” code on the C side, and that code tends to be sensitively dependent on the internal implementation details of the Lisp system.
- The Lisp system can automatically convert objects back and forth between the Lisp and foreign representations. This is convenient, but translation becomes prohibitively slow when large or complex data structures must be shared. This approach is supported by the SBCL FFI, and used automatically when passing integers and strings.
- The Lisp program can directly manipulate foreign objects through the use of extensions to the Lisp language.

SBCL, like CMUCL before it, relies primarily on the automatic conversion and direct manipulation approaches. The SB-ALIEN package provides a facility wherein foreign values of simple scalar types are automatically converted and complex types are directly manipulated in their foreign representation. Additionally the lower-level System Area Pointers (or SAPs) can be used where necessary to provide untyped access to foreign memory.

Any foreign objects that can’t automatically be converted into Lisp values are represented by objects of type `alien-value`. Since Lisp is a dynamically typed language, even foreign objects must have a runtime type; this type information is provided by encapsulating the raw pointer to the foreign data within an `alien-value` object.

The type language and operations on foreign types are intentionally similar to those of the C language.

9.2 Foreign Types

Alien types have a description language based on nested list structure. For example the C type

```
struct foo {
    int a;
    struct foo *b[100];
};
```

has the corresponding SBCL FFI type

```
(struct foo
  (a int)
  (b (array (* (struct foo)) 100)))
```

9.2.1 Defining Foreign Types

Types may be either named or anonymous. With structure and union types, the name is part of the type specifier, allowing recursively defined types such as:

```
(struct foo (a (* (struct foo))))
```

An anonymous structure or union type is specified by using the name `nil`. The `with-alien` macro defines a local scope which “captures” any named type definitions. Other types are not inherently named, but can be given named abbreviations using the `define-alien-type` macro.

9.2.2 Foreign Types and Lisp Types

The foreign types form a subsystem of the SBCL type system. An `alien` type specifier provides a way to use any foreign type as a Lisp type specifier. For example,

```
(typep foo '(alien (* int)))
```

can be used to determine whether `foo` is a pointer to a foreign `int`. `alien` type specifiers can be used in the same ways as ordinary Lisp type specifiers (like `string`.) Alien type declarations are subject to the same precise type checking as any other declaration. See Section 4.2.2 [Precise Type Checking], page 19.

Note that the type identifiers used in the foreign type system overlap with native Lisp type specifiers in some cases. For example, the type specifier `(alien single-float)` is identical to `single-float`, since foreign floats are automatically converted to Lisp floats. When `type-of` is called on an alien value that is not automatically converted to a Lisp value, then it will return an `alien` type specifier.

9.2.3 Foreign Type Specifiers

Note: All foreign type names are exported from the `sb-alien` package. Some foreign type names are also symbols in the `common-lisp` package, in which case they are reexported from the `sb-alien` package, so that e.g. it is legal to refer to `sb-alien:single-float`.

These are the basic foreign type specifiers:

- The foreign type specifier `(* foo)` describes a pointer to an object of type `foo`. A pointed-to type `foo` of `t` indicates a pointer to anything, similar to `void *` in ANSI C. A null alien pointer can be detected with the `sb-alien:null-alien` function.
- The foreign type specifier `(array foo &rest dimensions)` describes array of the specified `dimensions`, holding elements of type `foo`. Note that (unlike in C) `(* foo)` and `(array foo)` are considered to be different types when type checking is done. If equivalence of pointer and array types is desired, it may be explicitly coerced using `sb-alien:cast`.

Arrays are accessed using `sb-alien:deref`, passing the indices as additional arguments. Elements are stored in column-major order (as in C), so the first dimension determines only the size of the memory block, and not the layout of the higher dimensions. An array whose first dimension is variable may be specified by using `nil` as the first dimension. Fixed-size arrays can be allocated as array elements, structure slots or `sb-alien:with-alien` variables. Dynamic arrays can only be allocated using `sb-alien:make-alien`.

- The foreign type specifier `(sb-alien:struct name &rest fields)` describes a structure type with the specified `name` and `fields`. Fields are allocated at the same offsets used by the implementation's C compiler, as guessed by the SBCL internals. An optional `:alignment` keyword argument can be specified for each field to explicitly control the alignment of a field. If `name` is `nil` then the structure is anonymous.

If a named foreign `struct` specifier is passed to `define-alien-type` or `with-alien`, then this defines, respectively, a new global or local foreign structure type. If no `fields` are specified, then the fields are taken from the current (local or global) alien structure type definition of `name`.

- The foreign type specifier `(sb-alien:union name &rest fields)` is similar to `sb-alien:struct`, but describes a union type. All fields are allocated at the same offset, and the size of the union is the size of the largest field. The programmer must determine which field is active from context.
- The foreign type specifier `(sb-alien:enum name &rest specs)` describes an enumeration type that maps between integer values and symbols. If `name` is `nil`, then the type is anonymous. Each element of the `specs` list is either a Lisp symbol, or a list `(symbol value)`. `value` is an integer. If `value` is not supplied, then it defaults to one greater than the value for the preceding spec (or to zero if it is the first spec).
- The foreign type specifier `(sb-alien:signed &optional bits)` specifies a signed integer with the specified number of `bits` precision. The upper limit on integer precision is determined by the machine's word size. If `bits` is not specified, the maximum size will be used.
- The foreign type specifier `(integer &optional bits)` is equivalent to the corresponding type specifier using `sb-alien:signed` instead of `integer`.
- The foreign type specifier `(sb-alien:unsigned &optional bits)` is like corresponding type specifier using `sb-alien:signed` except that the variable is treated as an unsigned integer.

- The foreign type specifier (**boolean** &**optional bits**) is similar to an enumeration type, but maps from Lisp `nil` and `t` to C `0` and `1` respectively. *bits* determines the amount of storage allocated to hold the truth value.
- The foreign type specifier **single-float** describes a floating-point number in IEEE single-precision format.
- The foreign type specifier **double-float** describes a floating-point number in IEEE double-precision format.
- The foreign type specifier (**function result-type** &**rest arg-types**) describes a foreign function that takes arguments of the specified *arg-types* and returns a result of type *result-type*. Note that the only context where a foreign **function** type is directly specified is in the argument to **sb-alien:alien-funcall**. In all other contexts, foreign functions are represented by foreign function pointer types: **(* (function ...))**.
- The foreign type specifier **sb-alien:system-area-pointer** describes a pointer which is represented in Lisp as a **system-area-pointer** object. SBCL exports this type from **sb-alien** because CMUCL did, but tentatively (as of the first draft of this section of the manual, SBCL 0.7.6) it is deprecated, since it doesn't seem to be required by user code.
- The foreign type specifier **sb-alien:void** is used in function types to declare that no useful value is returned. Using **alien-funcall** to call a void foreign function will return zero values.
- The foreign type specifier (**sb-alien:c-string** &**key external-format element-type not-null**) is similar to **(* char)**, but is interpreted as a null-terminated string, and is automatically converted into a Lisp string when accessed; or if the pointer is C `NULL` or `0`, then accessing it gives Lisp `nil` unless **not-null** is true, in which case a type-error is signalled.

External format conversion is automatically done when Lisp strings are passed to foreign code, or when foreign strings are passed to Lisp code. If the type specifier has an explicit **external-format**, that external format will be used. Otherwise a default external format that has been determined at SBCL startup time based on the current locale settings will be used. For example, when the following alien routine is called, the Lisp string given as argument is converted to an **ebcdic** octet representation.

```
(define-alien-routine test int (str (c-string :external-format :ebcdic-us)))
```

Lisp strings of type **base-string** are stored with a trailing NUL termination, so no copying (either by the user or the implementation) is necessary when passing them to foreign code, assuming that the **external-format** and **element-type** of the **c-string** type are compatible with the internal representation of the string. For an SBCL built with Unicode support that means an **external-format** of **:ascii** and an **element-type** of **base-char**. Without Unicode support the **external-format** can also be **:iso-8859-1**, and the **element-type** can also be **character**. If the **external-format** or **element-type** is not compatible, or the string is a **(simple-array character *)**, this data is copied by the implementation as required.

Assigning a Lisp string to a **c-string** structure field or variable stores the contents of the string to the memory already pointed to by that variable. When a foreign object of type **(* char)** is assigned to a **c-string**, then the **c-string** pointer is assigned to. This allows **c-string** pointers to be initialized. For example:

```
(cl:in-package "CL-USER") ; which USEs package "SB-ALIEN"

(define-alien-type nil (struct foo (str c-string)))

(defun make-foo (str)
  (let ((my-foo (make-alien (struct foo))))
    (setf (slot my-foo 'str) (make-alien char (length str))
          (slot my-foo 'str) str)
    my-foo))
```

Storing Lisp `NIL` in a **c-string** writes C `NULL` to the variable.

- **sb-alien** also exports translations of these C type specifiers as foreign type specifiers: **char**, **short**, **int**, **long**, **unsigned-char**, **unsigned-short**, **unsigned-int**, **unsigned-long**, **float**, **double**, **size-t**, and **off-t**.

9.3 Operations On Foreign Values

This section describes how to read foreign values as Lisp values, how to coerce foreign values to different kinds of foreign values, and how to dynamically allocate and free foreign variables.

9.3.1 Accessing Foreign Values

sb-alien:deref *pointer-or-array* &rest *indices* [Function]

The **sb-alien:deref** function returns the value pointed to by a foreign pointer, or the value of a foreign array element. When dereferencing a pointer, an optional single index can be specified to give the equivalent of C pointer arithmetic; this index is scaled by the size of the type pointed to. When dereferencing an array, the number of indices must be the same as the number of dimensions in the array type. **deref** can be set with **setf** to assign a new value.

sb-alien:slot *struct-or-union* *slot-name* [Function]

The **sb-alien:slot** function extracts the value of the slot named *slot-name* from a foreign **struct** or **union**. If *struct-or-union* is a pointer to a structure or union, then it is automatically dereferenced. **sb-alien:slot** can be set with **setf** to assign a new value. Note that *slot-name* is evaluated, and need not be a compile-time constant (but only constant slot accesses are efficiently compiled).

9.3.1.1 Untyped memory

As noted at the beginning of the chapter, the System Area Pointer facilities allow untyped access to foreign memory. SAPs can be converted to and from the usual typed foreign values using **sap-alien** and **alien-sap** (described elsewhere), and also to and from integers - raw machine addresses. They should thus be used with caution; corrupting the Lisp heap or other memory with SAPs is trivial.

sb-sys:int-sap *machine-address* [Function]

Creates a SAP pointing at the virtual address *machine-address*.

sb-sys:sap-ref-32 *sap* *offset* [Function]

Access the value of the memory location at *offset* bytes from *sap*. This form may also be used with **setf** to alter the memory at that location.

sb-sys:sap= *sap1* *sap2* [Function]

Compare *sap1* and *sap2* for equality.

Similarly named functions exist for accessing other sizes of word, other comparisons, and other conversions. The reader is invited to use **apropos** and **describe** for more details

(**apropos** "sap" :sb-sys)

9.3.2 Coercing Foreign Values

sb-alien:addr *alien-expr* [Macro]

The **sb-alien:addr** macro returns a pointer to the location specified by *alien-expr*, which must be either a foreign variable, a use of **sb-alien:deref**, a use of **sb-alien:slot**, or a use of **sb-alien:extern-alien**.

sb-alien:cast *foreign-value* *new-type* [Macro]

The **sb-alien:cast** macro converts *foreign-value* to a new foreign value with the specified *new-type*. Both types, old and new, must be foreign pointer, array or function types. Note that the resulting Lisp foreign variable object is not **eq** to the argument, but it does refer to the same foreign data bits.

sb-alien:sap-alien *sap* *type* [Macro]

The **sb-alien:sap-alien** macro converts *sap* (a system area pointer) to a foreign value with the specified *type*. *type* is not evaluated.

The *type* must be some foreign pointer, array, or record type.

sb-alien:alien-sap *foreign-value* [Function]

The **sb-alien:alien-sap** function returns the SAP which points to *alien-value*'s data.

The *foreign-value* must be of some foreign pointer, array, or record type.

9.3.3 Foreign Dynamic Allocation

Lisp code can call the C standard library functions `malloc` and `free` to dynamically allocate and deallocate foreign variables. The Lisp code shares the same allocator with foreign C code, so it's OK for foreign code to call `free` on the result of Lisp `sb-alien:make-alien`, or for Lisp code to call `sb-alien:free-alien` on foreign objects allocated by C code.

`sb-alien:make-alien` *type &optional size* [Macro]

Allocate an alien of type `type` in foreign heap, and return an alien pointer to it. The allocated memory is not initialized, and may contain garbage. The memory is allocated using `malloc(3)`, so it can be passed to foreign functions which use `free(3)`, or released using `free-alien`.

For alien stack allocation, see macro `with-alien`.

The `type` argument is not evaluated. If `size` is supplied, how it is interpreted depends on `type`:

- When `type` is a foreign array type, an array of that type is allocated, and a pointer to it is returned. Note that you must use `deref` to first access the array through the pointer.

If supplied, `size` is used as the first dimension for the array.

- When `type` is any other foreign type, then an object for that type is allocated, and a pointer to it is returned. So `(make-alien int)` returns a `(* int)`.

If `size` is specified, then a block of that many objects is allocated, with the result pointing to the first one.

Examples:

```
(defvar *foo* (make-alien (array char 10)))
(type-of *foo*)           ; => (alien (* (array (signed 8) 10)))
(setf (deref (deref *foo*) 0) 10) ; => 10

(make-alien char 12)      ; => (alien (* (signed 8)))
```

`sb-alien:make-alien-string` *string &rest rest &key start end external-format* [Function]
null-terminate

Copy part of `string` delimited by `start` and `end` into freshly allocated foreign memory, freeable using `free(3)` or `free-alien`. Returns the allocated string as a `(* char)` alien, and the number of bytes allocated as secondary value.

The string is encoded using `external-format`. If `null-terminate` is true (the default), the alien string is terminated by an additional null byte.

`sb-alien:free-alien` *alien* [Function]

Dispose of the storage pointed to by `alien`. The `alien` must have been allocated by `make-alien`, `make-alien-string` or `malloc(3)`.

9.4 Foreign Variables

Both local (stack allocated) and external (C global) foreign variables are supported.

9.4.1 Local Foreign Variables

`sb-alien:with-alien` *var-definitions &body body* [Macro]

The `with-alien` macro establishes local foreign variables with the specified alien types and names. This form is analogous to defining a local variable in C: additional storage is allocated, and the initial value is copied. This form is less analogous to LET-allocated Lisp variables, since the variables can't be captured in closures: they live only for the dynamic extent of the body, and referring to them outside is a gruesome error.

The `var-definitions` argument is a list of variable definitions, each of the form

```
(name type &optional initial-value)
```

The names of the variables are established as symbol-macros; the bindings have lexical scope, and may be assigned with `setq` or `setf`.

The `with-alien` macro also establishes a new scope for named structures and unions. Any `type` specified for a variable may contain named structure or union types with the slots specified. Within the lexical scope of the binding specifiers and body, a locally defined foreign structure type `foo` can be referenced by its name using `(struct foo)`.

9.4.2 External Foreign Variables

External foreign names are strings, and Lisp names are symbols. When an external foreign value is represented using a Lisp variable, there must be a way to convert from one name syntax into the other. The macros `extern-alien`, `define-alien-variable` and `define-alien-routine` use this conversion heuristic:

- Alien names are converted to Lisp names by uppercasing and replacing underscores with hyphens.
- Conversely, Lisp names are converted to alien names by lowercasing and replacing hyphens with underscores.
- Both the Lisp symbol and alien string names may be separately specified by using a list of the form
(alien-string lisp-symbol)

`sb-alien:define-alien-variable` *name type* [Macro]

The `define-alien-variable` macro defines *name* as an external foreign variable of the specified foreign *type*. *name* and *type* are not evaluated. The Lisp name of the variable (see above) becomes a global alien variable. Global alien variables are effectively “global symbol macros”; a reference to the variable fetches the contents of the external variable. Similarly, setting the variable stores new contents – the new contents must be of the declared *type*. Someday, they may well be implemented using the ANSI `define-symbol-macro` mechanism, but as of SBCL 0.7.5, they are still implemented using an older more-or-less parallel mechanism inherited from CMUCL.

For example, to access a C-level counter *foo*, one could write

```
(define-alien-variable "foo" int)
;; Now it is possible to get the value of the C variable foo simply by
;; referencing that Lisp variable:
(print foo)
(setf foo 14)
(incf foo)
```

`sb-alien:get-errno` [Function]

Since in modern C libraries, the `errno` “variable” is typically no longer a variable, but some bizarre artificial construct which behaves superficially like a variable within a given thread, it can no longer reliably be accessed through the ordinary `define-alien-variable` mechanism. Instead, SBCL provides the operator `sb-alien:get-errno` to allow Lisp code to read it.

`sb-alien:extern-alien` *name type* [Macro]

The `extern-alien` macro returns an alien with the specified *type* which points to an externally defined value. *name* is not evaluated, and may be either a string or a symbol. *type* is an unevaluated alien type specifier.

9.5 Foreign Data Structure Examples

Now that we have alien types, operations and variables, we can manipulate foreign data structures. This C declaration

```
struct foo {
    int a;
    struct foo *b[100];
};
```

can be translated into the following alien type:

```
(define-alien-type nil
  (struct foo
    (a int)
    (b (array (* (struct foo)) 100))))
```

Once the `foo` alien type has been defined as above, the C expression

```
struct foo f;
f.b[7].a;
```

can be translated in this way:

```
(with-alien ((f (struct foo)))
  (slot (deref (slot f 'b) 7) 'a)
  ;;
```

```
;; Do something with f...
)
```

Or consider this example of an external C variable and some accesses:

```
struct c_struct {
    short x, y;
    char a, b;
    int z;
    c_struct *n;
};
extern struct c_struct *my_struct;
my_struct->x++;
my_struct->a = 5;
my_struct = my_struct->n;
```

which can be manipulated in Lisp like this:

```
(define-alien-type nil
  (struct c-struct
    (x short)
    (y short)
    (a char)
    (b char)
    (z int)
    (n (* c-struct))))
(define-alien-variable "my_struct" (* c-struct))
(incf (slot my-struct 'x))
(setf (slot my-struct 'a) 5)
(setq my-struct (slot my-struct 'n))
```

9.6 Loading Shared Object Files

Foreign object files can be loaded into the running Lisp process by calling `load-shared-object`.

`sb-alien:load-shared-object pathname &key dont-save` [Function]
Load a shared library / dynamic shared object file / similar foreign container specified by designated *pathname*, such as a .so on an elf platform.

Locating the shared object follows standard rules of the platform, consult the manual page for `dlopen(3)` for details. Typically paths specified by environment variables such as `LD_LIBRARY_PATH` are searched if the *pathname* has no directory, but on some systems (eg. Mac OS X) search may happen even if *pathname* is absolute. (On Windows `LoadLibrary` is used instead of `dlopen(3)`.)

On non-Windows platforms calling `load-shared-object` again with a *pathname* equal to the designated *pathname* of a previous call will replace the old definitions; if a symbol was previously referenced through the object and is not present in the reloaded version an error will be signalled. Reloading may not work as expected if user or library-code has called `dlopen(3)` on the same shared object or running on a system where `dlclose(3)` is a noop.

`load-shared-object` interacts with `sb-ext:save-lisp-and-die`:

1. If *dont-save* is true (default is `nil`), the shared object will be dropped when `save-lisp-and-die` is called -- otherwise shared objects are reloaded automatically when a saved core starts up. Specifying *dont-save* can be useful when the location of the shared object on startup is uncertain.
2. On most platforms references in compiled code to foreign symbols in shared objects (such as those generated by `define-alien-routine`) remain valid across `save-lisp-and-die`. On those platforms where this is not supported, a warning will be signalled when the core is saved -- this is orthogonal from *dont-save*.

`sb-alien:unload-shared-object pathname` [Function]
Unloads the shared object loaded earlier using the designated *pathname* with `load-shared-object`, to the degree supported on the platform.

Experimental.

9.7 Foreign Function Calls

The foreign function call interface allows a Lisp program to call many functions written in languages that use the C calling convention.

Lisp sets up various signal handling routines and other environment information when it first starts up, and expects these to be in place at all times. The C functions called by Lisp should not change the environment, especially the signal handlers: the signal handlers installed by Lisp typically have interesting flags set (e.g to request machine context information, or for signal delivery on an alternate stack) which the Lisp runtime relies on for correct operation. Precise details of how this works may change without notice between versions; the source, or the brain of a friendly SBCL developer, is the only documentation. Users of a Lisp built with the `:sb-thread` feature should also read the section about threads, Chapter 13 [Threading], page 96.

9.7.1 The `alien-funcall` Primitive

`sb-alien:alien-funcall` *alien-function* &rest *arguments* [Function]

The `alien-funcall` function is the foreign function call primitive: *alien-function* is called with the supplied *arguments* and its C return value is returned as a Lisp value. The *alien-function* is an arbitrary run-time expression; to refer to a constant function, use `extern-alien` or a value defined by `define-alien-routine`.

The type of *alien-function* must be `(alien (function ...))` or `(alien (* (function ...)))`. The function type is used to determine how to call the function (as though it was declared with a prototype.) The type need not be known at compile time, but only known-type calls are efficiently compiled. Limitations:

- Structure type return values are not implemented.
- Passing of structures by value is not implemented.

Here is an example which allocates a `(struct foo)`, calls a foreign function to initialize it, then returns a Lisp vector of all the `(* (struct foo))` objects filled in by the foreign call:

```
;; Allocate a foo on the stack.
(with-alien ((f (struct foo)))
  ;; Call some C function to fill in foo fields.
  (alien-funcall (extern-alien "mangle_foo" (function void (* foo)))
    (addr f))
  ;; Find how many foos to use by getting the A field.
  (let* ((num (slot f 'a))
        (result (make-array num)))
    ;; Get a pointer to the array so that we don't have to keep extracting it:
    (with-alien ((a (* (array (* (struct foo)) 100)) (addr (slot f 'b))))
      ;; Loop over the first N elements and stash them in the result vector.
      (dotimes (i num)
        (setf (svref result i) (deref (deref a) i)))
      ;; Voila.
      result)))
```

9.7.2 The `define-alien-routine` Macro

`sb-alien:define-alien-routine` *name* *result-type* &rest *arg-specifiers* [Macro]

The `define-alien-routine` macro is a convenience for automatically generating Lisp interfaces to simple foreign functions. The primary feature is the parameter style specification, which translates the C pass-by-reference idiom into additional return values.

name is usually a string external symbol, but may also be a symbol Lisp name or a list of the foreign name and the Lisp name. If only one name is specified, the other is automatically derived as for `extern-alien`. *result-type* is the alien type of the return value.

Each element of the *arg-specifiers* list specifies an argument to the foreign function, and is of the form

`(aname atype &optional style)`

aname is the symbol name of the argument to the constructed function (for documentation). *atype* is the alien type of corresponding foreign argument. The semantics of the actual call are the same as for

alien-funcall. *style* specifies how this argument should be handled at call and return time, and should be one of the following:

- **:in** specifies that the argument is passed by value. This is the default. **:in** arguments have no corresponding return value from the Lisp function.
- **:copy** is similar to **:in**, but the argument is copied to a pre-allocated object and a pointer to this object is passed to the foreign routine.
- **:out** specifies a pass-by-reference output value. The type of the argument must be a pointer to a fixed-sized object (such as an integer or pointer). **:out** and **:in-out** style cannot be used with pointers to arrays, records or functions. An object of the correct size is allocated on the stack, and its address is passed to the foreign function. When the function returns, the contents of this location are returned as one of the values of the Lisp function (and the location is automatically deallocated).
- **:in-out** is a combination of **:copy** and **:out**. The argument is copied to a pre-allocated object and a pointer to this object is passed to the foreign routine. On return, the contents of this location is returned as an additional value.

Note: Any efficiency-critical foreign interface function should be inline expanded, which can be done by preceding the **define-alien-routine** call with:

```
(declare (inline lisp-name))
```

In addition to avoiding the Lisp call overhead, this allows pointers, word-integers and floats to be passed using non-descriptor representations, avoiding consing.)

9.7.3 define-alien-routine Example

Consider the C function `cfoo` with the following calling convention:

```
void
cfoo (str, a, i)
    char *str;
    char *a; /* update */
    int *i; /* out */
{
    /* body of cfoo(...) */
}
```

This can be described by the following call to **define-alien-routine**:

```
(define-alien-routine "cfoo" void
  (str c-string)
  (a char :in-out)
  (i int :out))
```

The Lisp function `cfoo` will have two arguments (`str` and `a`) and two return values (`a` and `i`).

9.8 Calling Lisp From C

SBCL supports the calling of Lisp functions using the C calling convention. This is useful for both defining callbacks and for creating an interface for calling into Lisp as a shared library directly from C.

The **define-alien-callable** macro wraps Lisp code and creates a C foreign function which can be called with the C calling convention.

sb-alien:define-alien-callable *name result-type typed-lambda-list &body body* [Macro]
 Define an alien callable function in the alien callable namespace with result type **result-type** and with lambda list specifying the alien types of the arguments.

The **alien-callable-function** function returns the foreign callable value associated with any name defined by **define-alien-callable**, so that we can, for example, pass the callable value to C as a callback.

sb-alien:alien-callable-function *name* [Function]
 Return the alien callable function associated with **name**.

Note that the garbage collector moves objects, and won't be able to fix up any references in C variables. There are three mechanisms for coping with this:

1. The `sb-ext:purify` moves all live Lisp data into static or read-only areas such that it will never be moved (or freed) again in the life of the Lisp session
2. `sb-sys:with-pinned-objects` is a macro which arranges for some set of objects to be pinned in memory for the dynamic extent of its body forms. On ports which use the generational garbage collector (most, as of this writing) this affects exactly the specified objects. On other ports it is implemented by turning off GC for the duration (so could be said to have a whole-world granularity).
3. Disable GC, using the `without-gcing` macro.

9.8.1 Lisp as a Shared Library

SBCL supports the use of Lisp as a shared library that can be used by C programs using the `define-alien-callable` interface. See the `:callable-exports` keyword to `save-lisp-and-die` for how to save the Lisp image in a way that allows a C program to initialize the Lisp runtime and the exported symbols. When SBCL is built as a library, it exposes the symbol `initialize_lisp` which can be used in conjunction with a core initializing global symbols to foreign callables as function pointers and with object code allocating those symbols to initialize the runtime properly. The arguments to `initialize_lisp` are the same as the arguments to the main `sbcl` program.

Note: There is currently no way to run exit hooks or otherwise undo Lisp initialization gracefully from C.

9.9 Step-By-Step Example of the Foreign Function Interface

This section presents a complete example of an interface to a somewhat complicated C function.

Suppose you have the following C function which you want to be able to call from Lisp in the file `test.c`

```
struct c_struct
{
    int x;
    char *s;
};

struct c_struct *c_function (i, s, r, a)
{
    int i;
    char *s;
    struct c_struct *r;
    int a[10];

    int j;
    struct c_struct *r2;

    printf("i = %d\n", i);
    printf("s = %s\n", s);
    printf("r->x = %d\n", r->x);
    printf("r->s = %s\n", r->s);
    for (j = 0; j < 10; j++) printf("a[%d] = %d.\n", j, a[j]);
    r2 = (struct c_struct *) malloc (sizeof(struct c_struct));
    r2->x = i + 5;
    r2->s = "a C string";
    return(r2);
};
```

It is possible to call this C function from Lisp using the file `test.lisp` containing

```
(cl:defpackage "TEST-C-CALL" (:use "CL" "SB-ALIEN" "SB-C-CALL"))
(cl:in-package "TEST-C-CALL")

;;; Define the record C-STRUCT in Lisp.
(define-alien-type nil
  (struct c-struct
    (x int)
    (s c-string)))
```



```

;;; Define the Lisp function interface to the C routine. It returns a
;;; pointer to a record of type C-STRUCT. It accepts four parameters:
;;; I, an int; S, a pointer to a string; R, a pointer to a C-STRUCT
;;; record; and A, a pointer to the array of 10 ints.
;;;
;;; The INLINE declaration eliminates some efficiency notes about heap
;;; allocation of alien values.
(declare (inline c-function))
(define-alien-routine c-function
  (* (struct c-struct))
  (i int)
  (s c-string)
  (r (* (struct c-struct)))
  (a (array int 10)))

;;; a function which sets up the parameters to the C function and
;;; actually calls it
(defun call-cfun ()
  (with-alien ((ar (array int 10))
               (c-struct (struct c-struct)))
    (dotimes (i 10) ; Fill array.
      (setf (deref ar i) i))
    (setf (slot c-struct 'x) 20)
    (setf (slot c-struct 's) "a Lisp string")

    (with-alien ((res (* (struct c-struct))
                        (c-function 5 "another Lisp string" (addr c-struct) ar)))
      (format t "~&back from C function~%")
      (multiple-value-prog1
        (values (slot res 'x)
                (slot res 's))

        ; Deallocate result. (after we are done referring to it:
        ; "Pillage, *then* burn.")
        (free-alien res))))))

```

To execute the above example, it is necessary to compile the C routine, e.g.: `'cc -c test.c && ld -shared -o test.so test.o'` (In order to enable incremental loading with some linkers, you may need to say `'cc -G 0 -c test.c'`)

Once the C code has been compiled, you can start up Lisp and load it in: `'sbcl'`. Lisp should start up with its normal prompt.

Within Lisp, compile the Lisp file. (This step can be done separately. You don't have to recompile every time.) `'(compile-file "test.lisp")'`

Within Lisp, load the foreign object file to define the necessary symbols: `'(load-shared-object "test.so")'`.

Now you can load the compiled Lisp ("fasl") file into Lisp: `'(load "test.fasl")'` And once the Lisp file is loaded, you can call the Lisp routine that sets up the parameters and calls the C function: `'(test-c-call::call-cfun)'`

The C routine should print the following information to standard output:

```

i = 5
s = another Lisp string
r->x = 20
r->s = a Lisp string
a[0] = 0.
a[1] = 1.
a[2] = 2.
a[3] = 3.

```

```
a[4] = 4.  
a[5] = 5.  
a[6] = 6.  
a[7] = 7.  
a[8] = 8.  
a[9] = 9.
```

After return from the C function, the Lisp wrapper function should print the following output:

```
back from C function
```

And upon return from the Lisp wrapper function, before the next prompt is printed, the Lisp read-eval-print loop should print the following return values:

```
10  
"a C string"
```

10 Pathnames

10.1 Lisp Pathnames

There are many aspects of ANSI Common Lisp's pathname support which are implementation-defined and so need documentation.

10.1.1 Home Directory Specifiers

SBCL accepts the keyword `:home` and a list of the form `(:home "username")` as a directory component immediately following `:absolute`.

`:home` is represented in namestrings by `~/` and `(:home "username")` by `~username/` at the start of the namestring. Tilde-characters elsewhere in namestrings represent themselves.

Home directory specifiers are resolved to home directory of the current or specified user by `native-namestring`, which is used by the implementation to translate pathnames before passing them on to operating system specific routines.

Using `(:home "user")` form on Windows signals an error.

10.1.2 The SYS Logical Pathname Host

The logical pathname host named by "SYS" exists in SBCL. Its `logical-pathname-translations` may be set by the site or the user applicable to point to the locations of the system's sources; in particular, the core system's source files match the logical pathname `"SYS:SRC;**/*.*.*"`, and the contributed modules' source files match `"SYS:CONTRIB;**/*.*.*"`.

`sb-ext:set-sbcl-source-location pathname` [Function]

Initialize the `sys` logical host based on `pathname`, which should be the top-level directory of the `sbcl` sources. This will replace any existing translations for `"SYS:SRC;"`, `"SYS:CONTRIB;"`, and `"SYS:OUTPUT;"`. Other `"SYS:"` translations are preserved.

10.2 Native Filenames

In some circumstances, what is wanted is a Lisp pathname object which corresponds to a string produced by the Operating System. In this case, some of the default parsing rules are inappropriate: most filesystems do not have a native understanding of wild pathnames; such functionality is often provided by shells above the OS, often in mutually-incompatible ways.

To allow the user to deal with this, the following functions are provided: `parse-native-namestring` and `native-pathname` return the closest equivalent Lisp pathname to a given string (appropriate for the Operating System), while `native-namestring` converts a non-wild pathname designator to the equivalent native namestring, if possible. Some Lisp pathname concepts (such as the `:back` directory component) have no direct equivalents in most Operating Systems; the behaviour of `native-namestring` is unspecified if an inappropriate pathname designator is passed to it. Additionally, note that conversion from pathname to native filename and back to pathname should not be expected to preserve equivalence under `equal`.

`sb-ext:parse-native-namestring thing &optional host defaults &key start end` [Function]
junk-allowed as-directory

Convert `thing` into a pathname, using the native conventions appropriate for the pathname host `host`, or if not specified the host of `defaults`. If `thing` is a string, the parse is bounded by `start` and `end`, and error behaviour is controlled by `junk-allowed`, as with `parse-namestring`. For file systems whose native conventions allow directories to be indicated as files, if `as-directory` is true, return a pathname denoting `thing` as a directory.

`sb-ext:native-pathname pathspec` [Function]

Convert `pathspec` (a pathname designator) into a pathname, assuming the operating system native pathname conventions.

`sb-ext:native-namestring pathname &key as-file` [Function]

Construct the full native (name)string form of `pathname`. For file systems whose native conventions allow directories to be indicated as files, if `as-file` is true and the name, type, and version components of `pathname` are all `nil` or `:unspecific`, construct a string that names the directory according to the file system's syntax for files.

Because some file systems permit the names of directories to be expressed in multiple ways, it is occasionally necessary to parse a native file name “as a directory name” or to produce a native file name that names a directory “as a file”. For these cases, `parse-native-namestring` accepts the keyword argument `as-directory` to force a filename to parse as a directory, and `native-namestring` accepts the keyword argument `as-file` to force a pathname to unparse as a file. For example,

```
; On Unix, the directory "/tmp/" can be denoted by "/tmp/" or "/tmp".
; Under the default rules for native filenames, these parse and
; unparse differently.
(defvar *p*)
(setf *p* (parse-native-namestring "/tmp/")) => #P"/tmp/"
(pathname-name *p*) => NIL
(pathname-directory *p*) => (:ABSOLUTE "tmp")
(native-namestring *p*) => "/tmp/"

(setf *p* (parse-native-namestring "/tmp")) => #P"/tmp"
(pathname-name *p*) => "tmp"
(pathname-directory *p*) => (:ABSOLUTE)
(native-namestring *p*) => "/tmp"

; A non-NIL AS-DIRECTORY argument to PARSE-NATIVE-NAMESTRING forces
; both the second string to parse the way the first does.
(setf *p* (parse-native-namestring "/tmp"
                                   nil *default-pathname-defaults*
                                   :as-directory t)) => #P"/tmp/"

(pathname-name *p*) => NIL
(pathname-directory *p*) => (:ABSOLUTE "tmp")

; A non-NIL AS-FILE argument to NATIVE-NAMESTRING forces the pathname
; parsed from the first string to unparse as the second string.
(setf *p* (parse-native-namestring "/tmp/")) => #P"/tmp/"
(native-namestring *p* :as-file t) => "/tmp"
```

11 Streams

Streams which read or write Lisp character data from or to the outside world – files, sockets or other external entities – require the specification of a conversion between the external, binary data and the Lisp characters. In ANSI Common Lisp, this is done by specifying the `:external-format` argument when the stream is created. The major information required is an *encoding*, specified by a keyword naming that encoding; however, it is also possible to specify refinements to that encoding as additional options to the external format designator.

In addition, SBCL supports various extensions of ANSI Common Lisp streams:

Bivalent Streams

A type of stream that can read and write both `character` and `(unsigned-byte 8)` values.

Gray Streams

User-overloadable CLOS classes whose instances can be used as Lisp streams (e.g. passed as the first argument to `format`).

Simple Streams

The bundled contrib module *sb-simple-streams* implements a subset of the Franz Allegro simple-streams proposal.

11.1 Stream External Formats

The function `stream-external-format` returns the canonical name of the external format (See Chapter 8 [External Formats], page 67) used by the stream for character-based input and/or output.

When constructing file streams, for example using `open` or `with-open-file`, the external format to use is specified via the `:external-format` argument which accepts an external format designator (See Section 8.2 [External Format Designators], page 67).

11.2 Bivalent Streams

A *bivalent stream* can be used to read and write both `character` and `(unsigned-byte 8)` values. A bivalent stream is created by calling `open` with the argument `:element-type :default`. On such a stream, both binary and character data can be read and written with the usual input and output functions.

Streams are *not* created bivalent by default for performance reasons. Bivalent streams are incompatible with `fast-read-char`, an internal optimization in SBCL's stream machinery that bulk-converts octets to characters and implements a fast path through `read-char`.

11.3 Gray Streams

The Gray Streams interface is a widely supported extension that provides for definition of CLOS-extensible stream classes. Gray stream classes are implemented by adding methods to generic functions analogous to Common Lisp's standard I/O functions. Instances of Gray stream classes may be used with any I/O operation where a non-Gray stream can, provided that all required methods have been implemented suitably.

11.3.1 Gray Streams classes

The defined Gray Stream classes are these:

`sb-gray:fundamental-stream` [Class]

Class precedence list: `fundamental-stream`, `standard-object`, `stream`, `t`

Base class for all Gray streams.

`sb-gray:fundamental-input-stream` [Class]

Class precedence list: `fundamental-input-stream`, `fundamental-stream`, `standard-object`, `stream`, `t`

Superclass of all Gray input streams.

The function `input-stream-p` will return true of any generalized instance of `fundamental-input-stream`.

sb-gray:fundamental-output-stream [Class]
 Class precedence list: **fundamental-output-stream**, **fundamental-stream**, **standard-object**, **stream**, **t**
 Superclass of all Gray output streams.

The function **output-stream-p** will return true of any generalized instance of **fundamental-output-stream**.

sb-gray:fundamental-binary-stream [Class]
 Class precedence list: **fundamental-binary-stream**, **fundamental-stream**, **standard-object**, **stream**, **t**
 Superclass of all Gray streams whose element-type is a subtype of unsigned-byte or signed-byte.

Note that instantiable subclasses of **fundamental-binary-stream** should provide (or inherit) an applicable method for the generic function **stream-element-type**.

sb-gray:fundamental-character-stream [Class]
 Class precedence list: **fundamental-character-stream**, **fundamental-stream**, **standard-object**, **stream**, **t**
 Superclass of all Gray streams whose element-type is a subtype of character.

sb-gray:fundamental-binary-input-stream [Class]
 Class precedence list: **fundamental-binary-input-stream**, **fundamental-input-stream**, **fundamental-binary-stream**, **fundamental-stream**, **standard-object**, **stream**, **t**
 Superclass of all Gray input streams whose element-type is a subtype of unsigned-byte or signed-byte.

sb-gray:fundamental-binary-output-stream [Class]
 Class precedence list: **fundamental-binary-output-stream**, **fundamental-output-stream**, **fundamental-binary-stream**, **fundamental-stream**, **standard-object**, **stream**, **t**
 Superclass of all Gray output streams whose element-type is a subtype of unsigned-byte or signed-byte.

sb-gray:fundamental-character-input-stream [Class]
 Class precedence list: **fundamental-character-input-stream**, **fundamental-input-stream**, **fundamental-character-stream**, **fundamental-stream**, **standard-object**, **stream**, **t**
 Superclass of all Gray input streams whose element-type is a subtype of character.

sb-gray:fundamental-character-output-stream [Class]
 Class precedence list: **fundamental-character-output-stream**, **fundamental-output-stream**, **fundamental-character-stream**, **fundamental-stream**, **standard-object**, **stream**, **t**
 Superclass of all Gray output streams whose element-type is a subtype of character.

11.3.2 Methods common to all streams

These generic functions can be specialized on any generalized instance of **fundamental-stream**.

c1:stream-element-type stream [Generic Function]
 Return a type specifier for the kind of object returned by the **stream**. The class **fundamental-character-stream** provides a default method which returns **character**.

c1:close stream &key abort [Generic Function]
 Close the given **stream**. No more I/O may be performed, but inquiries may still be made. If **:abort** is true, an attempt is made to clean up the side effects of having created the stream.

sb-gray:stream-file-position stream &optional position-spec [Generic Function]
 Used by **file-position**. Returns or changes the current position within **stream**.

11.3.3 Input stream methods

These generic functions may be specialized on any generalized instance of **fundamental-input-stream**.

sb-gray:stream-clear-input stream [Generic Function]
 This is like **c1:clear-input**, but for Gray streams, returning **nil**. The default method does nothing.

sb-gray:stream-read-sequence stream seq &optional start end [Generic Function]
 This is like **c1:read-sequence**, but for Gray streams.

11.3.4 Character input stream methods

These generic functions are used to implement subclasses of `fundamental-input-stream`:

`sb-gray:stream-peek-char stream` [Generic Function]

This is used to implement `peek-char`; this corresponds to `peek-type` of `nil`. It returns either a character or `:eof`. The default method calls `stream-read-char` and `stream-unread-char`.

`sb-gray:stream-read-char-no-hang stream` [Generic Function]

This is used to implement `read-char-no-hang`. It returns either a character, or `nil` if no input is currently available, or `:eof` if end-of-file is reached. The default method provided by `fundamental-character-input-stream` simply calls `stream-read-char`; this is sufficient for file streams, but interactive streams should define their own method.

`sb-gray:stream-read-char stream` [Generic Function]

Read one character from the stream. Return either a character object, or the symbol `:eof` if the stream is at end-of-file. Every subclass of `fundamental-character-input-stream` must define a method for this function.

`sb-gray:stream-read-line stream` [Generic Function]

This is used by `read-line`. A string is returned as the first value. The second value is true if the string was terminated by end-of-file instead of the end of a line. The default method uses repeated calls to `stream-read-char`.

`sb-gray:stream-listen stream` [Generic Function]

This is used by `listen`. It returns true or false. The default method uses `stream-read-char-no-hang` and `stream-unread-char`. Most streams should define their own method since it will usually be trivial and will always be more efficient than the default method.

`sb-gray:stream-unread-char stream character` [Generic Function]

Undo the last call to `stream-read-char`, as in `unread-char`. Return `nil`. Every subclass of `fundamental-character-input-stream` must define a method for this function.

11.3.5 Output stream methods

These generic functions are used to implement subclasses of `fundamental-output-stream`:

`sb-gray:stream-clear-output stream` [Generic Function]

This is like `cl:clear-output`, but for Gray streams: clear the given output `stream`. The default method does nothing.

`sb-gray:stream-finish-output stream` [Generic Function]

Attempts to ensure that all output sent to the Stream has reached its destination, and only then returns false. Implements `finish-output`. The default method does nothing.

`sb-gray:stream-force-output stream` [Generic Function]

Attempts to force any buffered output to be sent. Implements `force-output`. The default method does nothing.

`sb-gray:stream-write-sequence stream seq &optional start end` [Generic Function]

This is like `cl:write-sequence`, but for Gray streams.

11.3.6 Character output stream methods

These generic functions are used to implement subclasses of `fundamental-character-output-stream`:

`sb-gray:stream-advance-to-column stream column` [Generic Function]

Write enough blank space so that the next character will be written at the specified column. Returns true if the operation is successful, or `nil` if it is not supported for this stream. This is intended for use by `pprint` and `format ~T`. The default method uses `stream-line-column` and repeated calls to `stream-write-char` with a `#space` character; it returns `nil` if `stream-line-column` returns `nil`.

`sb-gray:stream-fresh-line stream` [Generic Function]

Outputs a new line to the Stream if it is not positioned at the beginning of a line. Returns `t` if it output a new line, `nil` otherwise. Used by `fresh-line`. The default method uses `stream-start-line-p` and `stream-terpri`.

sb-gray:stream-line-column *stream* [Generic Function]

Return the column number where the next character will be written, or `nil` if that is not meaningful for this stream. The first column on a line is numbered 0. This function is used in the implementation of `pprint` and the `format ~T` directive. For every character output stream class that is defined, a method must be defined for this function, although it is permissible for it to always return `nil`.

sb-gray:stream-line-length *stream* [Generic Function]

Return the stream line length or `nil`.

sb-gray:stream-start-line-p *stream* [Generic Function]

Is `stream` known to be positioned at the beginning of a line? It is permissible for an implementation to always return `nil`. This is used in the implementation of `fresh-line`. Note that while a value of 0 from `stream-line-column` also indicates the beginning of a line, there are cases where `stream-start-line-p` can be meaningfully implemented although `stream-line-column` can't be. For example, for a window using variable-width characters, the column number isn't very meaningful, but the beginning of the line does have a clear meaning. The default method for `stream-start-line-p` on class `fundamental-character-output-stream` uses `stream-line-column`, so if that is defined to return `nil`, then a method should be provided for either `stream-start-line-p` or `stream-fresh-line`.

sb-gray:stream-terpri *stream* [Generic Function]

Writes an end of line, as for `terpri`. Returns `nil`. The default method does (`stream-write-char stream #newline`).

sb-gray:stream-write-char *stream character* [Generic Function]

Write `character` to `stream` and return `character`. Every subclass of `fundamental-character-output-stream` must have a method defined for this function.

sb-gray:stream-write-string *stream string &optional start end* [Generic Function]

This is used by `write-string`. It writes the string to the stream, optionally delimited by `start` and `end`, which default to 0 and `nil`. The string argument is returned. The default method provided by `fundamental-character-output-stream` uses repeated calls to `stream-write-char`.

11.3.7 Binary stream methods

The following generic functions are available for subclasses of `fundamental-binary-stream`:

sb-gray:stream-read-byte *stream* [Generic Function]

Used by `read-byte`; returns either an integer, or the symbol `:eof` if the stream is at end-of-file.

sb-gray:stream-write-byte *stream integer* [Generic Function]

Implements `write-byte`; writes the integer to the stream and returns the integer as the result.

11.3.8 Gray Streams examples

Below are two classes of stream that can be conveniently defined as wrappers for Common Lisp streams. These are meant to serve as examples of minimal implementations of the protocols that must be followed when defining Gray streams. Realistic uses of the Gray Streams API would implement the various methods that can do I/O in batches, such as `stream-read-line`, `stream-write-string`, `stream-read-sequence`, and `stream-write-sequence`.

11.3.8.1 Character counting input stream

It is occasionally handy for programs that process input files to count the number of characters and lines seen so far, and the number of characters seen on the current line, so that useful messages may be reported in case of parsing errors, etc. Here is a character input stream class that keeps track of these counts. Note that all character input streams must implement `stream-read-char` and `stream-unread-char`.

```
(defclass wrapped-stream (fundamental-stream)
  ((stream :initarg :stream :reader stream-of)))
(defmethod stream-element-type ((stream wrapped-stream))
  (stream-element-type (stream-of stream)))
(defmethod close ((stream wrapped-stream) &key abort)
  (close (stream-of stream) :abort abort))
```



```

(defclass wrapped-character-input-stream
  (wrapped-stream fundamental-character-input-stream)
  ())
(defmethod stream-read-char ((stream wrapped-character-input-stream))
  (read-char (stream-of stream) nil :eof))
(defmethod stream-unread-char ((stream wrapped-character-input-stream)
                               char)
  (unread-char char (stream-of stream)))
(defclass counting-character-input-stream
  (wrapped-character-input-stream)
  ((char-count :initform 1 :accessor char-count-of)
   (line-count :initform 1 :accessor line-count-of)
   (col-count :initform 1 :accessor col-count-of)
   (prev-col-count :initform 1 :accessor prev-col-count-of)))
(defmethod stream-read-char ((stream counting-character-input-stream))
  (with-accessors ((inner-stream stream-of) (chars char-count-of)
                  (lines line-count-of) (cols col-count-of)
                  (prev prev-col-count-of)) stream
    (let ((char (call-next-method)))
      (cond ((eql char :eof)
              :eof)
            ((char= char #\Newline)
             (incf lines)
             (incf chars)
             (setf prev cols)
             (setf cols 1)
             char)
            (t
             (incf chars)
             (incf cols)
             char)))))
(defmethod stream-unread-char ((stream counting-character-input-stream)
                               char)
  (with-accessors ((inner-stream stream-of) (chars char-count-of)
                  (lines line-count-of) (cols col-count-of)
                  (prev prev-col-count-of)) stream
    (cond ((char= char #\Newline)
            (decf lines)
            (decf chars)
            (setf cols prev))
          (t
           (decf chars)
           (decf cols)
           char))
    (call-next-method)))

```

The default methods for `stream-read-char-no-hang`, `stream-peek-char`, `stream-listen`, `stream-clear-input`, `stream-read-line`, and `stream-read-sequence` should be sufficient (though the last two will probably be slower than methods that forwarded directly).

Here's a sample use of this class:

```

(with-input-from-string (input "1 2
3 :foo ")
  (let ((counted-stream (make-instance 'counting-character-input-stream
                                     :stream input)))
    (loop for thing = (read counted-stream) while thing
          unless (numberp thing) do
            (error "Non-number ~S (line ~D, column ~D)" thing
                  (line-count-of counted-stream)
                  (- (col-count-of counted-stream)
                     (length (format nil "~S" thing))))
          end
          do (print thing))))
1
2
3
Non-number :FOO (line 2, column 5)
[Condition of type SIMPLE-ERROR]

```

11.3.8.2 Output prefixing character stream

One use for a wrapped output stream might be to prefix each line of text with a timestamp, e.g., for a logging stream. Here's a simple stream that does this, though without any fancy line-wrapping. Note that all character output stream classes must implement `stream-write-char` and `stream-line-column`.

```

(defclass wrapped-stream (fundamental-stream)
  ((stream :initarg :stream :reader stream-of)))
(defmethod stream-element-type ((stream wrapped-stream))
  (stream-element-type (stream-of stream)))
(defmethod close ((stream wrapped-stream) &key abort)
  (close (stream-of stream) :abort abort))
(defclass wrapped-character-output-stream
  (wrapped-stream fundamental-character-output-stream)
  ((col-index :initform 0 :accessor col-index-of)))
(defmethod stream-line-column ((stream wrapped-character-output-stream))
  (col-index-of stream))
(defmethod stream-write-char ((stream wrapped-character-output-stream)
                             char)
  (with-accessors ((inner-stream stream-of) (cols col-index-of)) stream
    (write-char char inner-stream)
    (if (char= char #\Newline)
        (setf cols 0)
        (incf cols))))
(defclass prefixed-character-output-stream
  (wrapped-character-output-stream)
  ((prefix :initarg :prefix :reader prefix-of)))
(defgeneric write-prefix (prefix stream)
  (:method ((prefix string) stream) (write-string prefix stream))
  (:method ((prefix function) stream) (funcall prefix stream)))
(defmethod stream-write-char ((stream prefixed-character-output-stream)
                             char)
  (with-accessors ((inner-stream stream-of) (cols col-index-of)
                  (prefix prefix-of)) stream
    (when (zerop cols)
      (write-prefix prefix inner-stream))
    (call-next-method)))

```

As with the example input stream, this implements only the minimal protocol. A production implementation should also provide methods for at least `stream-write-line`, `stream-write-sequence`.

And here's a sample use of this class:

```

(flet ((format-timestamp (stream)
      (apply #'format stream "[~2@*~2,' D:~1@*~2,'OD:~0@*~2,'OD] "
              (multiple-value-list (get-decoded-time)))))
  (let ((output (make-instance 'prefixed-character-output-stream
                              :stream *standard-output*
                              :prefix #'format-timestamp)))
    (loop for string in '("abc" "def" "ghi") do
      (write-line string output)
      (sleep 1))))
[ 0:30:05] abc
[ 0:30:06] def
[ 0:30:07] ghi
NIL

```

11.4 Simple Streams

Simple streams are an extensible streams protocol that avoids some problems with Gray streams.

Documentation about simple streams is available at:

<http://www.franz.com/support/documentation/6.2/doc/streams.htm>

The implementation should be considered Alpha-quality; the basic framework is there, but many classes are just stubs at the moment.

See `SYS:CONTRIB;SB-SIMPLE-STREAMS;SIMPLE-STREAM-TEST.LISP` for things that should work.

Known differences to the ACL behaviour:

- `open` not return a simple-stream by default. This can be adjusted; see `default-open-class` in the file `cl.lisp`
- `write-vector` is unimplemented.

12 Package Locks

None of the following sections apply to SBCL built without package locking support.

warning: The interface described here is experimental: incompatible changes in future SBCL releases are possible, even expected: the concept of “implementation packages” and the associated operators may be renamed; more operations (such as naming restarts or catch tags) may be added to the list of operations violating package locks.

12.1 Package Lock Concepts

12.1.1 Package Locking Overview

Package locks protect against unintentional modifications of a package: they provide similar protection to user packages as is mandated to `common-lisp` package by the ANSI specification. They are not, and should not be used as, a security measure.

Newly created packages are by default unlocked (see the `:lock` option to `defpackage`).

The package `common-lisp` and SBCL internal implementation packages are locked by default, including `sb-ext`.

It may be beneficial to lock `common-lisp-user` as well, to ensure that various libraries don’t pollute it without asking, but this is not currently done by default.

12.1.2 Implementation Packages

Each package has a list of associated implementation packages. A locked package, and the symbols whose home package it is, can be modified without violating package locks only when `*package*` is bound to one of the implementation packages of the locked package.

Unless explicitly altered by `defpackage`, `sb-ext:add-implementation-package`, or `sb-ext:remove-implementation-package` each package is its own (only) implementation package.

12.1.3 Package Lock Violations

12.1.3.1 Lexical Bindings and Declarations

Lexical bindings or declarations that violate package locks cause a compile-time warning, and a runtime `program-error` when the form that violates package locks would be executed.

A complete listing of operators affected by this is: `let`, `let*`, `flet`, `labels`, `macrolet`, and `symbol-macrolet`, `declare`.

Package locks affecting both lexical bindings and declarations can be disabled locally with `sb-ext:disable-package-locks` declaration, and re-enabled with `sb-ext:enable-package-locks` declaration.

Example:

```
(in-package :locked)

(defun foo () ...)

(defmacro with-foo (&body body)
  `(locally (declare (disable-package-locks locked:foo))
    (flet ((foo () ...))
      (declare (enable-package-locks locked:foo)) ; re-enable for body
      ,@body)))
```

12.1.3.2 Other Operations

If a non-lexical operation violates a package lock, a `continuable` error that is of a subtype of `sb-ext:package-lock-violation` (subtype of `package-error`) is signalled when the operation is attempted.

Additional restarts may be established for `continuable` package lock violations for interactive use.

The actual type of the error depends on circumstances that caused the violation: operations on packages signal errors of type `sb-ext:package-locked-error`, and operations on symbols signal errors of type `sb-ext:symbol-package-locked-error`.

12.1.4 Package Locks in Compiled Code

12.1.4.1 Interned Symbols

If file-compiled code contains interned symbols, then loading that code into an image without the said symbols will not cause a package lock violation, even if the packages in question are locked.

12.1.4.2 Other Limitations on Compiled Code

With the exception of interned symbols, behaviour is unspecified if package locks affecting compiled code are not the same during loading of the code or execution.

Specifically, code compiled with packages unlocked may or may not fail to signal package-lock-violations even if the packages are locked at runtime, and code compiled with packages locked may or may not signal spurious package-lock-violations at runtime even if the packages are unlocked.

In practice all this means that package-locks have a negligible performance penalty in compiled code as long as they are not violated.

12.1.5 Operations Violating Package Locks

12.1.5.1 Operations on Packages

The following actions cause a package lock violation if the package operated on is locked, and `*package*` is not an implementation package of that package, and the action would cause a change in the state of the package (so e.g. exporting already external symbols is never a violation). Package lock violations caused by these operations signal errors of type `sb-ext:package-locked-error`.

1. Shadowing a symbol in a package.
2. Importing a symbol to a package.
3. Uninterning a symbol from a package.
4. Exporting a symbol from a package.
5. Unexporting a symbol from a package.
6. Changing the packages used by a package.
7. Renaming a package.
8. Deleting a package.
9. Adding a new package local nickname to a package.
10. Removing an existing package local nickname to a package.

12.1.5.2 Operations on Symbols

Following actions cause a package lock violation if the home package of the symbol operated on is locked, and `*package*` is not an implementation package of that package. Package lock violations caused by these action signal errors of type `sb-ext:symbol-package-locked-error`.

These actions cause only one package lock violation per lexically apparent violated package.

Example:

```
;;; Packages FOO and BAR are locked.
;;;
;;; Two lexically apparent violated packages: exactly two
;;; package-locked-errors will be signalled.

(defclass foo:point ()
  ((x :accessor bar:x)
   (y :accessor bar:y)))
```

1. Binding or altering its value lexically or dynamically, or establishing it as a symbol-macro.
Exceptions:
 - If the symbol is not defined as a constant, global symbol-macro or a global dynamic variable, it may be lexically bound or established as a local symbol macro.
 - If the symbol is defined as a global dynamic variable, it may be assigned or bound.
2. Defining, undefining, or binding it, or its setf name as a function.
Exceptions:
 - If the symbol is not defined as a function, macro, or special operator it and its setf name may be lexically bound as a function.
3. Defining, undefining, or binding it as a macro or compiler macro.
Exceptions:
 - If the symbol is not defined as a function, macro, or special operator it may be lexically bound as a macro.
4. Defining it as a type specifier or structure.
5. Defining it as a declaration with a declaration proclamation.
6. Declaring or proclaiming it special.
7. Declaring or proclaiming its type or ftype.
Exceptions:
 - If the symbol may be lexically bound, the type of that binding may be declared.
 - If the symbol may be lexically bound as a function, the ftype of that binding may be declared.
8. Defining a setf expander for it.
9. Defining it as a method combination type.
10. Using it as the class-name argument to setf of find-class.
11. Defining it as a hash table test using `sb-ext:define-hash-table-test`.

12.2 Package Lock Dictionary

`sb-ext:disable-package-locks`

[Declaration]

Syntax: (`sb-ext:disable-package-locks` symbol*)

Disables package locks affecting the named symbols during compilation in the lexical scope of the declaration. Disabling locks on symbols whose home package is unlocked, or disabling an already disabled lock, has no effect.

`sb-ext:enable-package-locks`

[Declaration]

Syntax: (`sb-ext:enable-package-locks` symbol*)

Re-enables package locks affecting the named symbols during compilation in the lexical scope of the declaration. Enabling locks that were not first disabled with `sb-ext:disable-package-locks` declaration, or enabling locks that are already enabled has no effect.

`sb-ext:package-lock-violation`

[Condition]

Class precedence list: `package-lock-violation`, `package-error`, `error`, `serious-condition`, `condition`, `t`

Subtype of `cl:package-error`. A subtype of this error is signalled when a package-lock is violated.

`sb-ext:package-locked-error`

[Condition]

Class precedence list: `package-locked-error`, `package-lock-violation`, `package-error`, `error`, `serious-condition`, `condition`, `t`

Subtype of `sb-ext:package-lock-violation`. An error of this type is signalled when an operation on a package violates a package lock.

`sb-ext:symbol-package-locked-error`

[Condition]

Class precedence list: `symbol-package-locked-error`, `package-lock-violation`, `package-error`, `error`, `serious-condition`, `condition`, `t`

Subtype of `sb-ext:package-lock-violation`. An error of this type is signalled when an operation on a symbol violates a package lock. The symbol that caused the violation is accessed by the function `sb-ext:package-locked-error-symbol`.

`sb-ext:package-locked-error-symbol` *symbol-package-locked-error* [Function]
Returns the symbol that caused the `symbol-package-locked-error` condition.

`sb-ext:package-locked-p` *package* [Function]
Returns `t` when `package` is locked, `nil` otherwise. Signals an error if `package` doesn't designate a valid package.

`sb-ext:lock-package` *package* [Function]
Locks `package` and returns `t`. Has no effect if `package` was already locked. Signals an error if `package` is not a valid package designator

`sb-ext:unlock-package` *package* [Function]
Unlocks `package` and returns `t`. Has no effect if `package` was already unlocked. Signals an error if `package` is not a valid package designator.

`sb-ext:package-implemented-by-list` *package* [Function]
Returns a list containing the implementation packages of `package`. Signals an error if `package` is not a valid package designator.

`sb-ext:package-implements-list` *package* [Function]
Returns the packages that `package` is an implementation package of. Signals an error if `package` is not a valid package designator.

`sb-ext:add-implementation-package` *packages-to-add &optional package* [Function]
Adds `packages-to-add` as implementation packages of `package`. Signals an error if `package` or any of the `packages-to-add` is not a valid package designator.

`sb-ext:remove-implementation-package` *packages-to-remove &optional package* [Function]
Removes `packages-to-remove` from the implementation packages of `package`. Signals an error if `package` or any of the `packages-to-remove` is not a valid package designator.

`sb-ext:without-package-locks` *&body body* [Macro]
Ignores all runtime package lock violations during the execution of `body`. `Body` can begin with declarations.

`sb-ext:with-unlocked-packages` (*&rest packages*) *&body forms* [Macro]
Unlocks `packages` for the dynamic scope of the `body`. Signals an error if any of `packages` is not a valid package designator.

`cl:defpackage` *name* *[[option]]** \Rightarrow *package* [Macro]
Options are extended to include the following:

- `:lock` *boolean*

If the argument to `:lock` is `t`, the package is initially locked. If `:lock` is not provided it defaults to `nil`.

- `:implement` *package-designator**

The package is added as an implementation package to the packages named. If `:implement` is not provided, it defaults to the package itself.

Example:

```
(defpackage "FOO" (:export "BAR") (:lock t) (:implement))
(defpackage "FOO-INT" (:use "FOO") (:implement "FOO" "FOO-INT"))
```

;;; is equivalent to

```
(defpackage "FOO") (:export "BAR")
(lock-package "FOO")
(remove-implementation-package "FOO" "FOO")
(defpackage "FOO-INT" (:use "BAR"))
(add-implementation-package "FOO-INT" "FOO")
```

13 Threading

SBCL supports a fairly low-level threading interface that maps onto the host operating system's concept of threads or lightweight processes. This means that threads may take advantage of hardware multiprocessing on machines that have more than one CPU, but it does not allow Lisp control of the scheduler. This is found in the SB-THREAD package.

Threads are part of the default build on x86[-64]/ARM64 Linux and Windows.

They are also supported on: x86[-64] Darwin (Mac OS X), x86[-64] FreeBSD, x86 SunOS (Solaris), PPC Linux, ARM64 Linux, RISC-V Linux. On these platforms threads must be explicitly enabled at build-time, see `INSTALL` for directions.

13.1 Threading basics

```
(make-thread (lambda () (write-line "Hello, world")))
```

13.1.1 Thread Objects

`sb-thread:thread` [Structure]

Class precedence list: `thread`, `structure-object`, `t`

Thread type. Do not rely on threads being structs as it may change in future versions.

`sb-thread:*current-thread*` [Variable]

Bound in each thread to the thread itself.

`sb-thread:list-all-threads` [Function]

Return a list of the live threads. Note that the return value is potentially stale even before the function returns, as new threads may be created and old ones may exit at any time.

`sb-thread:thread-alive-p thread` [Function]

Return `t` if `thread` is still alive. Note that the return value is potentially stale even before the function returns, as the thread may exit at any time.

`sb-thread:thread-name instance` [Function]

Name of the thread. Can be assigned to using `setf`. A thread name must be a simple-string (not necessarily unique) or `nil`.

`sb-thread:main-thread-p &optional thread` [Function]

True if `thread`, defaulting to current thread, is the main thread of the process.

`sb-thread:main-thread` [Function]

Returns the main thread of the process.

13.1.2 Making, Returning From, Joining, and Yielding Threads

`sb-thread:make-thread function &key name arguments` [Function]

Create a new thread of `name` that runs `function` with the argument list designator provided (defaults to no argument). Thread exits when the function returns. The return values of `function` are kept around and can be retrieved by `join-thread`.

Invoking the initial `abort` restart established by `make-thread` terminates the thread.

See also: `return-from-thread`, `abort-thread`.

`sb-thread:return-from-thread values-form &key allow-exit` [Macro]

Unwinds from and terminates the current thread, with values from `values-form` as the results visible to `join-thread`.

If current thread is the main thread of the process (see `main-thread-p`), signals an error unless `allow-exit` is true, as terminating the main thread would terminate the entire process. If `allow-exit` is true, returning from the main thread is equivalent to calling `sb-ext:exit` with `:code 0` and `:abort nil`.

See also: `abort-thread` and `sb-ext:exit`.

sb-thread:abort-thread *&key allow-exit* [Function]

Unwinds from and terminates the current thread abnormally, causing **join-thread** on current thread to signal an error unless a default-value is provided.

If current thread is the main thread of the process (see **main-thread-p**), signals an error unless **allow-exit** is true, as terminating the main thread would terminate the entire process. If **allow-exit** is true, aborting the main thread is equivalent to calling **sb-ext:exit** code 1 and **:abort nil**.

Invoking the initial **abort** restart established by **make-thread** is equivalent to calling **abort-thread** in other than main threads. However, whereas **abort** restart may be rebound, **abort-thread** always unwinds the entire thread. (Behaviour of the initial **abort** restart for main thread depends on the **:toplevel** argument to **sb-ext:save-lisp-and-die**.)

See also: **return-from-thread** and **sb-ext:exit**.

sb-thread:join-thread *thread &key default timeout* [Function]

Suspend current thread until **thread** exits. Return the result values of the thread function.

If **thread** does not exit within **timeout** seconds and **default** is supplied, return two values: 1) **default** 2) **:timeout**. If **default** is not supplied, signal a **join-thread-error** with **join-thread-problem** equal to **:timeout**.

If **thread** does not exit normally (i.e. aborted) and **default** is supplied, return two values: 1) **default** 2) **:abort**. If **default** is not supplied, signal a **join-thread-error** with **join-thread-problem** equal to **:abort**.

If **thread** is the current thread, signal a **join-thread-error** with **join-thread-problem** equal to **:self-join**.

Trying to join the main thread causes **join-thread** to block until **timeout** occurs or the process exits: when the main thread exits, the entire process exits.

Users should not rely on the ability to join a chosen **thread** from more than one other thread simultaneously. Future changes to **join-thread** may directly call the underlying thread library, and not all threading implementations consider such usage to be well-defined.

note: Return convention in case of a timeout is experimental and subject to change.

sb-thread:thread-yield [Function]

Yield the processor to other threads.

13.1.3 Asynchronous Operations

sb-thread:interrupt-thread *thread function* [Function]

Interrupt **thread** and make it run **function**.

The interrupt is asynchronous, and can occur anywhere with the exception of sections protected using **sb-sys:without-interrupts**.

function is called with interrupts disabled, under **sb-sys:allow-with-interrupts**. Since functions such as **grab-mutex** may try to enable interrupts internally, in most cases **function** should either enter **sb-sys:with-interrupts** to allow nested interrupts, or **sb-sys:without-interrupts** to prevent them completely.

When a thread receives multiple interrupts, they are executed in the order they were sent -- first in, first out.

This means that a great degree of care is required to use **interrupt-thread** safely and sanely in a production environment. The general recommendation is to limit uses of **interrupt-thread** for interactive debugging, banning it entirely from production environments -- it is simply exceedingly hard to use correctly.

With those caveats in mind, what you need to know when using it:

- If calling **function** causes a non-local transfer of control (ie. an unwind), all normal cleanup forms will be executed.

However, if the interrupt occurs during cleanup forms of an **unwind-protect**, it is just as if that had happened due to a regular **go**, **throw**, or **return-from**: the interrupted cleanup form and those following it in the same **unwind-protect** do not get executed.

sbcl tries to keep its own internals **asynch-unwind-safe**, but this is frankly an unreasonable expectation for third party libraries, especially given that **asynch-unwind-safety** does not compose: a function calling only **asynch-unwind-safe** function isn't automatically **asynch-unwind-safe**.

This means that in order for an asynch unwind to be safe, the entire callstack at the point of interruption needs to be asynch-unwind-safe.

- In addition to asynch-unwind-safety you must consider the issue of reentrancy. `interrupt-thread` can cause function that are never normally called recursively to be re-entered during their dynamic contour, which may cause them to misbehave. (Consider binding of special variables, values of global variables, etc.)

Taken together, these two restrict the "safe" things to do using `interrupt-thread` to a fairly minimal set. One useful one -- exclusively for interactive development use is using it to force entry to debugger to inspect the state of a thread:

```
(interrupt-thread thread #'break)
```

Short version: be careful out there.

`sb-thread:terminate-thread thread` [Function]

Terminate the thread identified by `thread`, by interrupting it and causing it to call `sb-ext:abort-thread` with `:allow-exit t`.

The unwind caused by `terminate-thread` is asynchronous, meaning that eg. thread executing

```
(let (foo)
  (unwind-protect
    (progn
      (setf foo (get-foo))
      (work-on-foo foo))
    (when foo
      ;; An interrupt occurring inside the cleanup clause
      ;; will cause cleanups from the current UNWIND-PROTECT
      ;; to be dropped.
      (release-foo foo))))
```

might miss calling `release-foo` despite `get-foo` having returned true if the interrupt occurs inside the cleanup clause, eg. during execution of `release-foo`.

Thus, in order to write an asynch unwind safe `unwind-protect` you need to use `without-interrupts`:

```
(let (foo)
  (sb-sys:without-interrupts
    (unwind-protect
      (progn
        (setf foo (sb-sys:allow-with-interrupts
                    (get-foo)))
        (sb-sys:with-local-interrupts
          (work-on-foo foo)))
      (when foo
        (release-foo foo)))))
```

Since most libraries using `unwind-protect` do not do this, you should never assume that unknown code can safely be terminated using `terminate-thread`.

13.1.4 Miscellaneous Operations

`sb-thread:symbol-value-in-thread symbol thread &optional errorp` [Function]

Return the local value of `symbol` in `thread`, and a secondary value of `t` on success.

If the value cannot be retrieved (because the thread has exited or because it has no local binding for `name`) and `errorp` is true signals an error of type `symbol-value-in-thread-error`; if `errorp` is false returns a primary value of `nil`, and a secondary value of `nil`.

Can also be used with `setf` to change the thread-local value of `symbol`.

`symbol-value-in-thread` is primarily intended as a debugging tool, and not as a mechanism for inter-thread communication.

13.1.5 Error Conditions

`sb-thread:thread-error` [Condition]

Class precedence list: `thread-error`, `error`, `serious-condition`, `condition`, `t`

Conditions of type `thread-error` are signalled when thread operations fail. The offending thread is initialized by the `:thread` initialization argument and read by the function `thread-error-thread`.

sb-thread:thread-error-thread *condition* [Function]

Return the offending thread that the `thread-error` pertains to.

sb-thread:interrupt-thread-error [Condition]

Class precedence list: `interrupt-thread-error`, `thread-error`, `error`, `serious-condition`, `condition`, `t`

Signalled when interrupting a thread fails because the thread has already exited. The offending thread can be accessed using `thread-error-thread`.

sb-thread:join-thread-error [Condition]

Class precedence list: `join-thread-error`, `thread-error`, `error`, `serious-condition`, `condition`, `t`

Signalled when joining a thread fails due to abnormal exit of the thread to be joined. The offending thread can be accessed using `thread-error-thread`.

13.2 Special Variables

The interaction of special variables with multiple threads is mostly as one would expect, with behaviour very similar to other implementations.

- global special values are visible across all threads;
- bindings (e.g. using `LET`) are local to the thread;
- threads do not inherit dynamic bindings from the parent thread

The last point means that

```
(defparameter *x* 0)
(let ((*x* 1))
  (sb-thread:make-thread (lambda () (print *x*)))))
```

prints 0 and not 1 as of 0.9.6.

13.3 Atomic Operations

Following atomic operations are particularly useful for implementing lockless algorithms.

sb-ext:atomic-decf *place &optional diff* [Macro]

Atomically decrements `place` by `diff`, and returns the value of `place` before the decrement.

`place` must access one of the following:

- a `defstruct` slot with declared type `(unsigned-byte 64)` or `aref` of a `(simple-array (unsigned-byte 64) (*))` The type `sb-ext:word` can be used for these purposes.
- `car` or `cdr` (respectively `first` or `rest`) of a `cons`.
- a variable defined using `defglobal` with a proclaimed type of `fixnum`.

Macroexpansion is performed on `place` before expanding `atomic-decf`.

Decrementing is done using modular arithmetic, which is well-defined over two different domains:

- For structures and arrays, the operation accepts and produces an `(unsigned-byte 64)`, and `diff` must be of type `(signed-byte 64)`. `atomic-decf` of `#x0` by one results in `#xFFFFFFFFFFFFFFFF` being stored in `place`.
- For other places, the domain is `fixnum`, and `diff` must be a `fixnum`. `atomic-decf` of `#x-4000000000000000` by one results in `#x3FFFFFFFFFFFFFFF` being stored in `place`.

`diff` defaults to 1.

experimental: Interface subject to change.

sb-ext:atomic-incf *place &optional diff* [Macro]

Atomically increments `place` by `diff`, and returns the value of `place` before the increment.

`place` must access one of the following:

- a `defstruct` slot with declared type `(unsigned-byte 64)` or `aref` of a `(simple-array (unsigned-byte 64) (*))` The type `sb-ext:word` can be used for these purposes.

- `car` or `cdr` (respectively `first` or `rest`) of a `cons`.
- a variable defined using `defglobal` with a proclaimed type of `fixnum`.

Macroexpansion is performed on `place` before expanding `atomic-incf`.

Incrementing is done using modular arithmetic, which is well-defined over two different domains:

- For structures and arrays, the operation accepts and produces an (`unsigned-byte 64`), and `diff` must be of type (`signed-byte 64`). `atomic-incf` of `#xFFFFFFFFFFFFFFFF` by one results in `#x0` being stored in `place`.
- For other places, the domain is `fixnum`, and `diff` must be a `fixnum`. `atomic-incf` of `#x3FFFFFFFFFFFFFFFFF` by one results in `#x-4000000000000000` being stored in `place`.

`diff` defaults to 1.

`experimental`: Interface subject to change.

`sb-ext:atomic-pop place` [Macro]

Like `pop`, but atomic. `place` may be read multiple times before the operation completes -- the write does not occur until such time that no other thread modified `place` between the read and the write.

Works on all CASable places.

`sb-ext:atomic-push obj place` [Macro]

Like `push`, but atomic. `place` may be read multiple times before the operation completes -- the write does not occur until such time that no other thread modified `place` between the read and the write.

Works on all CASable places.

`sb-ext:atomic-update place update-fn &rest arguments` [Macro]

Updates `place` atomically to the value returned by calling function designated by `update-fn` with `arguments` and the previous value of `place`.

`place` may be read and `update-fn` evaluated and called multiple times before the update succeeds: atomicity in this context means that the value of `place` did not change between the time it was read, and the time it was replaced with the computed value.

`place` can be any place supported by `sb-ext:compare-and-swap`.

Examples:

```
;;; Conses T to the head of FOO-LIST.
(defstruct foo list)
(defvar *foo* (make-foo))
(atomic-update (foo-list *foo*) #'cons t)

(let ((x (cons :count 0)))
  (mapc #'sb-thread:join-thread
    (loop repeat 1000
      collect (sb-thread:make-thread
        (lambda ()
          (loop repeat 1000
            do (atomic-update (cdr x) #'1+)
              (sleep 0.00001))))))
    ;; Guaranteed to be (:COUNT . 1000000) -- if you replace
    ;; atomic update with (INCF (CDR X)) above, the result becomes
    ;; unpredictable.
    x)
```

`sb-ext:compare-and-swap place old new` [Macro]

Atomically stores `new` in `place` if `old` matches the current value of `place`. Two values are considered to match if they are `eq`. Returns the previous value of `place`: if the returned value is `eq` to `old`, the swap was carried out.

`place` must be an CAS-able place. Built-in CAS-able places are accessor forms whose `car` is one of the following:

`car`, `cdr`, `first`, `rest`, `svref`, `symbol-plist`, `symbol-value`, `svref`, `slot-value` `sb-mop:standard-instance-access`, `sb-mop:funcallable-standard-instance-access`,

or the name of a `defstruct` created accessor for a slot whose storage type is not raw. (Refer to the the "Efficiency" chapter of the manual for the list of raw slot types. Future extensions to this macro may allow it to work on some raw slot types.)

In case of `slot-value`, if the slot is unbound, `slot-unbound` is called unless `old` is `eq` to `sb-pcl:+slot-unbound+` in which case `sb-pcl:+slot-unbound+` is returned and `new` is assigned to the slot. Additionally, the results are unspecified if there is an applicable method on either `sb-mop:slot-value-using-class`, (`setf sb-mop:slot-value-using-class`), or `sb-mop:slot-boundp-using-class`.

Additionally, the `place` can be anything for which a CAS-function has been defined. (See `sb-ext:cas` for more information.)

CAS Protocol

Our `compare-and-swap` is user-extensible by defining functions named (CAS place), allowing users to add CAS support to new places.

`sb-ext:cas place old new`

[Macro]

Synonym for `compare-and-swap`.

Additionally `defun`, `defgeneric`, `defmethod`, `flet`, and `labels` can be also used to define CAS-functions analogously to SETF-functions:

```
(defvar *foo* nil)

(defun (cas foo) (old new)
  (cas (symbol-value '*foo*) old new))
```

First argument of a `cas` function is the expected old value, and the second argument of is the new value. Note that the system provides no automatic atomicity for `cas` functions, nor can it verify that they are atomic: it is up to the implementor of a `cas` function to ensure its atomicity.

`experimental`: Interface subject to change.

13.4 Mutex Support

Mutexes are used for controlling access to a shared resource. One thread is allowed to hold the mutex, others which attempt to take it will be made to wait until it's free. Threads are woken in the order that they go to sleep.

```
(defpackage :demo (:use "CL" "SB-THREAD" "SB-EXT"))

(in-package :demo)

(defvar *a-mutex* (make-mutex :name "my lock"))

(defun thread-fn ()
  (format t "Thread ~A running ~%" *current-thread*)
  (with-mutex (*a-mutex*)
    (format t "Thread ~A got the lock~%" *current-thread*)
    (sleep (random 5)))
  (format t "Thread ~A dropped lock, dying now~%" *current-thread*))

(make-thread #'thread-fn)
(make-thread #'thread-fn)
```

`sb-thread:mutex`

[Structure]

Class precedence list: `mutex`, `structure-object`, `t`

Mutex type.

`sb-thread:with-mutex (mutex &key wait-p timeout value) &body body`

[Macro]

Acquire `mutex` for the dynamic scope of `body`. If `wait-p` is true (the default), and the `mutex` is not immediately available, sleep until it is available.

If `timeout` is given, it specifies a relative timeout, in seconds, on how long the system should try to acquire the lock in the contested case.

If the mutex isn't acquired successfully due to either `wait-p` or `timeout`, the body is not executed, and `with-mutex` returns `nil`.

Otherwise body is executed with the mutex held by current thread, and `with-mutex` returns the values of body.

Historically `with-mutex` also accepted a `value` argument, which when provided was used as the new owner of the mutex instead of the current thread. This is no longer supported: if `value` is provided, it must be either `nil` or the current thread.

sb-thread:with-recursive-lock (*mutex &key wait-p timeout*) *&body body* [Macro]

Acquire `mutex` for the dynamic scope of `body`.

If `wait-p` is true (the default), and the `mutex` is not immediately available or held by the current thread, sleep until it is available.

If `timeout` is given, it specifies a relative timeout, in seconds, on how long the system should try to acquire the lock in the contested case.

If the mutex isn't acquired successfully due to either `wait-p` or `timeout`, the body is not executed, and `with-recursive-lock` returns `nil`.

Otherwise body is executed with the mutex held by current thread, and `with-recursive-lock` returns the values of body.

Unlike `with-mutex`, which signals an error on attempt to re-acquire an already held mutex, `with-recursive-lock` allows recursive lock attempts to succeed.

sb-thread:make-mutex *&key name* [Function]

Create a mutex.

sb-thread:mutex-name *instance* [Function]

The name of the mutex. Settable.

sb-thread:mutex-owner *mutex* [Function]

Current owner of the mutex, `nil` if the mutex is free. Naturally, this is racy by design (another thread may acquire the mutex after this function returns), it is intended for informative purposes. For testing whether the current thread is holding a mutex see `holding-mutex-p`.

sb-thread:mutex-value *mutex* [Function]

Current owner of the mutex, `nil` if the mutex is free. May return a stale value, use `mutex-owner` instead.

sb-thread:grab-mutex *mutex &key waitp timeout* [Function]

Acquire `mutex` for the current thread. If `waitp` is true (the default) and the mutex is not immediately available, sleep until it is available.

If `timeout` is given, it specifies a relative timeout, in seconds, on how long `grab-mutex` should try to acquire the lock in the contested case.

If `grab-mutex` returns `t`, the lock acquisition was successful. In case of `waitp` being `nil`, or an expired timeout, `grab-mutex` may also return `nil` which denotes that `grab-mutex` did -not- acquire the lock.

Notes:

- `grab-mutex` is not interrupt safe. The correct way to call it is:
(`without-interrupts ... (allow-with-interrupts (grab-mutex ...)) ...`)
`without-interrupts` is necessary to avoid an interrupt unwinding the call while the mutex is in an inconsistent state while `allow-with-interrupts` allows the call to be interrupted from sleep.
- (`grab-mutex <mutex> :timeout 0.0`) differs from (`grab-mutex <mutex> :waitp nil`) in that the former may signal a `deadline-timeout` if the global deadline was due already on entering `grab-mutex`.
The exact interplay of `grab-mutex` and deadlines are reserved to change in future versions.
- It is recommended that you use `with-mutex` instead of calling `grab-mutex` directly.

sb-thread:release-mutex *mutex &key if-not-owner* [Function]

Release `mutex` by setting it to `nil`. Wake up threads waiting for this mutex.

`release-mutex` is not interrupt safe: interrupts should be disabled around calls to it.

If the current thread is not the owner of the mutex then it silently returns without doing anything (if `if-not-owner` is `:punt`), signals a `warning` (if `if-not-owner` is `:warn`), or releases the mutex anyway (if `if-not-owner` is `:force`).

13.5 Semaphores

Semaphores are among other things useful for keeping track of a countable resource, e.g. messages in a queue, and sleep when the resource is exhausted.

- sb-thread:semaphore** [Structure]
 Class precedence list: `semaphore`, `structure-object`, `t`
 Semaphore type. The fact that a `semaphore` is a `structure-object` should be considered an implementation detail, and may change in the future.
- sb-thread:make-semaphore** *&key name count* [Function]
 Create a semaphore with the supplied `count` and `name`.
- sb-thread:signal-semaphore** *semaphore &optional n* [Function]
 Increment the count of `semaphore` by `n`. If there are threads waiting on this semaphore, then `n` of them is woken up.
- sb-thread:wait-on-semaphore** *semaphore &key n timeout notification* [Function]
 Decrement the count of `semaphore` by `n` if the count would not be negative.
 Else blocks until the semaphore can be decremented. Returns the new count of `semaphore` on success.
 If `timeout` is given, it is the maximum number of seconds to wait. If the count cannot be decremented in that time, returns `nil` without decrementing the count.
 If `notification` is given, it must be a `semaphore-notification` object whose `semaphore-notification-status` is `nil`. If `wait-on-semaphore` succeeds and decrements the count, the status is set to `t`.
- sb-thread:try-semaphore** *semaphore &optional n notification* [Function]
 Try to decrement the count of `semaphore` by `n`. If the count were to become negative, punt and return `nil`, otherwise return the new count of `semaphore`.
 If `notification` is given it must be a semaphore notification object with `semaphore-notification-status` of `nil`. If the count is decremented, the status is set to `t`.
- sb-thread:semaphore-count** *instance* [Function]
 Returns the current count of the semaphore `instance`.
- sb-thread:semaphore-name** *semaphore* [Function]
 The name of the semaphore `instance`. Settable.
- sb-thread:semaphore-notification** [Structure]
 Class precedence list: `semaphore-notification`, `structure-object`, `t`
 Semaphore notification object. Can be passed to `wait-on-semaphore` and `try-semaphore` as the `:notification` argument. Consequences are undefined if multiple threads are using the same notification object in parallel.
- sb-thread:make-semaphore-notification** [Function]
 Constructor for semaphore-notification objects. `semaphore-notification-status` is initially `nil`.
- sb-thread:semaphore-notification-status** *semaphore-notification* [Function]
 Returns `t` if a `wait-on-semaphore` or `try-semaphore` using `semaphore-notification` has succeeded since the notification object was created or cleared.
- sb-thread:clear-semaphore-notification** *semaphore-notification* [Function]
 Resets the `semaphore-notification` object for use with another call to `wait-on-semaphore` or `try-semaphore`.

13.6 Waitqueue/condition variables

These are based on the POSIX condition variable design, hence the annoyingly CL-conflicting name. For use when you want to check a condition and sleep until it's true. For example: you have a shared queue, a writer process checking "queue is empty" and one or more readers that need to know when "queue is not empty". It sounds simple, but is astonishingly easy to deadlock if another process runs when you weren't expecting it to.

There are three components:

- the condition itself (not represented in code)
- the condition variable (a.k.a. waitqueue) which proxies for it
- a lock to hold while testing the condition

Important stuff to be aware of:

- when calling condition-wait, you must hold the mutex. condition-wait will drop the mutex while it waits, and obtain it again before returning for whatever reason;
- likewise, you must be holding the mutex around calls to condition-notify;
- a process may return from condition-wait in several circumstances: it is not guaranteed that the underlying condition has become true. You must check that the resource is ready for whatever you want to do to it.

```
(defvar *buffer-queue* (make-waitqueue))
(defvar *buffer-lock* (make-mutex :name "buffer lock"))

(defvar *buffer* (list nil))

(defun reader ()
  (with-mutex (*buffer-lock*)
    (loop
      (condition-wait *buffer-queue* *buffer-lock*)
      (loop
        (unless *buffer* (return))
        (let ((head (car *buffer*)))
          (setf *buffer* (cdr *buffer*))
          (format t "reader ~A woke, read ~A~%"
                  *current-thread* head)))))))

(defun writer ()
  (loop
    (sleep (random 5))
    (with-mutex (*buffer-lock*)
      (let ((el (intern
                  (string (code-char
                          (+ (char-code #\A) (random 26)))))))
        (setf *buffer* (cons el *buffer*))
        (condition-notify *buffer-queue*)))))

(make-thread #'writer)
(make-thread #'reader)
(make-thread #'reader)
```

`sb-thread:waitqueue` [Structure]

Class precedence list: waitqueue, structure-object, t
Waitqueue type.

`sb-thread:make-waitqueue &key name` [Function]

Create a waitqueue.

`sb-thread:waitqueue-name instance` [Function]

The name of the waitqueue. Settable.

`sb-thread:condition-wait queue mutex &key timeout` [Function]

Atomically release `mutex` and start waiting on `queue` until another thread wakes us up using either `condition-notify` or `condition-broadcast` on `queue`, at which point we re-acquire `mutex` and return `t`.

Spurious wakeups are possible.

If `timeout` is given, it is the maximum number of seconds to wait, including both waiting for the wakeup and the time to re-acquire `mutex`. When neither a wakeup nor a re-acquisition occurs within the given time, returns `nil` without re-acquiring `mutex`.

If `condition-wait` unwinds, it may do so with or without `mutex` being held.

Important: Since `condition-wait` may return without `condition-notify` or `condition-broadcast` having occurred, the correct way to write code that uses `condition-wait` is to loop around the call, checking the associated data:

```
(defvar *data* nil)
(defvar *queue* (make-waitqueue))
(defvar *lock* (make-mutex))

;; Consumer
(defun pop-data (&optional timeout)
  (with-mutex (*lock*)
    (loop until *data*
      do (or (condition-wait *queue* *lock* :timeout timeout)
            ;; Lock not held, must unwind without touching *data*.
            (return-from pop-data nil)))
    (pop *data*)))

;; Producer
(defun push-data (data)
  (with-mutex (*lock*)
    (push data *data*)
    (condition-notify *queue*)))
```

`sb-thread:condition-notify queue &optional n` [Function]
 Notify `n` threads waiting on `queue`.

important: The same mutex that is used in the corresponding `condition-wait` must be held by this thread during this call.

`sb-thread:condition-broadcast queue` [Function]
 Notify all threads waiting on `queue`.

important: The same mutex that is used in the corresponding `condition-wait` must be held by this thread during this call.

13.7 Barriers

These are based on the Linux kernel barrier design, which is in turn based on the Alpha CPU memory model. They are presently implemented for x86, x86-64, PPC, ARM64, and RISC-V systems, and behave as compiler barriers on all other CPUs.

In addition to explicit use of the `sb-thread:barrier` macro, the following functions and macros also serve as `:memory` barriers:

- `sb-ext:atomic-decf`, `sb-ext:atomic-incf`, `sb-ext:atomic-push`, and `sb-ext:atomic-pop`.
- `sb-ext:compare-and-swap`.
- `sb-thread:grab-mutex`, `sb-thread:release-mutex`, `sb-thread:with-mutex` and `sb-thread:with-recursive-lock`.
- `sb-thread:signal-semaphore`, `sb-thread:try-semaphore` and `sb-thread:wait-on-semaphore`.
- `sb-thread:condition-wait`, `sb-thread:condition-notify` and `sb-thread:condition-broadcast`.

`sb-thread:barrier (kind) &body forms` [Macro]
 Insert a barrier in the code stream, preventing some sort of reordering.
`kind` should be one of:

```
:compiler  Prevent the compiler from reordering memory access across the barrier.
:memory    Prevent the cpu from reordering any memory access across the barrier.
:read      Prevent the cpu from reordering any read access across the barrier.
:write     Prevent the cpu from reordering any write access across the barrier.
```

:data-dependency

Prevent the cpu from reordering dependent memory reads across the barrier (requiring reads before the barrier to complete before any reads after the barrier that depend on them). This is a weaker form of the **:read** barrier.

forms is an implicit **progn**, evaluated before the barrier. **barrier** returns the values of the last form in **forms**.

The file "memory-barriers.txt" in the Linux kernel documentation is highly recommended reading for anyone programming at this level.

13.8 Sessions/Debugging

If the user has multiple views onto the same Lisp image (for example, using multiple terminals, or a windowing system, or network access) they are typically set up as multiple *sessions* such that each view has its own collection of foreground/background/stopped threads. A thread which wishes to create a new session can use **sb-thread:with-new-session** to remove itself from the current session (which it shares with its parent and siblings) and create a fresh one. # See also **sb-thread:make-listener-thread**.

Within a single session, threads arbitrate between themselves for the user's attention. A thread may be in one of three notional states: foreground, background, or stopped. When a background process attempts to print a repl prompt or to enter the debugger, it will stop and print a message saying that it has stopped. The user at his leisure may switch to that thread to find out what it needs. If a background thread enters the debugger, selecting any restart will put it back into the background before it resumes. Arbitration for the input stream is managed by calls to **sb-thread:get-foreground** (which may block) and **sb-thread:release-foreground**.

13.9 Foreign threads

Direct calls to **pthread_create** (instead of **MAKE-THREAD**) create threads that SBCL is not aware of, these are called foreign threads. Currently, it is not possible to run Lisp code in such threads. This means that the Lisp side signal handlers cannot work. The best solution is to start foreign threads with signals blocked, but since third party libraries may create threads, it is not always feasible to do so. As a workaround, upon receiving a signal in a foreign thread, SBCL changes the thread's sigmask to block all signals that it wants to handle and resends the signal to the current process which should land in a thread that does not block it, that is, a Lisp thread.

The resignalling trick cannot work for synchronously triggered signals (SIGSEGV and co), take care not to trigger any. Resignalling for synchronously triggered signals in foreign threads is subject to **--lose-on-corruption**, see Section 3.3.1 [Runtime Options], page 12.

13.10 Implementation (Linux x86/x86-64)

Threading is implemented using pthreads and some Linux specific bits like futexes.

On x86 the per-thread local bindings for special variables is achieved using the %fs segment register to point to a per-thread storage area. This may cause interesting results if you link to foreign code that expects threading or creates new threads, and the thread library in question uses %fs in an incompatible way. On x86-64 the r12 register has a similar role.

Queues require the **sys_futex()** system call to be available: this is the reason for the NPTL requirement. We test at runtime that this system call exists.

Garbage collection is done with the existing Conservative Generational GC. Allocation is done in small (typically 8k) regions: each thread has its own region so this involves no stopping. However, when a region fills, a lock must be obtained while another is allocated, and when a collection is required, all processes are stopped. This is achieved by sending them signals, which may make for interesting behaviour if they are interrupted in system calls. The streams interface is believed to handle the required system call restarting correctly, but this may be a consideration when making other blocking calls e.g. from foreign library code.

Large amounts of the SBCL library have not been inspected for thread-safety. Some of the obviously unsafe areas have large locks around them, so compilation and fast loading, for example, cannot be parallelized. Work is ongoing in this area.

A new thread by default is created in the same POSIX process group and session as the thread it was created by. This has an impact on keyboard interrupt handling: pressing your terminal's intr key (typically

Control-C) will interrupt all processes in the foreground process group, including Lisp threads that SBCL considers to be notionally ‘background’. This is undesirable, so background threads are set to ignore the SIGINT signal.

`sb-thread:make-listener-thread` in addition to creating a new Lisp session makes a new POSIX session, so that pressing *Control-C* in one window will not interrupt another listener - this has been found to be embarrassing.

14 Timers

SBCL supports a system-wide event scheduler implemented on top of `setitimer` that also works with threads but does not require a separate scheduler thread.

The following example schedules a timer that writes “Hello, world” after two seconds.

```
(schedule-timer (make-timer (lambda ()
                             (write-line "Hello, world")
                             (force-output)))
                2)
```

It should be noted that writing timer functions requires special care, as the dynamic environment in which they run is unpredictable: dynamic variable bindings, locks held, etc, all depend on whatever code was running when the timer fired. The following example should serve as a cautionary tale:

```
(defvar *foo* nil)

(defun show-foo ()
  (format t "~&foo=~S~%" *foo*)
  (force-output t))

(defun demo ()
  (schedule-timer (make-timer #'show-foo) 0.5)
  (schedule-timer (make-timer #'show-foo) 1.5)
  (let ((*foo* t))
    (sleep 1.0))
  (let ((*foo* :surprise!))
    (sleep 2.0)))
```

14.1 Timer Dictionary

`sb-ext:timer`

[Structure]

Class precedence list: `timer`, `structure-object`, `t`

Timer type. Do not rely on timers being structs as it may change in future versions.

`sb-ext:make-timer` *function &key name thread*

[Function]

Create a timer that runs *function* when triggered.

If a *thread* is supplied, *function* is run in that thread. If *thread* is `t`, a new thread is created for *function* each time the timer is triggered. If *thread* is `nil`, *function* is run in an unspecified thread.

When *thread* is not `t`, `interrupt-thread` is used to run *function* and the ordering guarantees of `interrupt-thread` apply. In that case, *function* runs with interrupts disabled but `with-interrupts` is allowed.

`sb-ext:timer-name` *timer*

[Function]

Return the name of *timer*.

`sb-ext:timer-scheduled-p` *timer &key delta*

[Function]

See if *timer* will still need to be triggered after *delta* seconds from now. For timers with a repeat interval it returns true.

`sb-ext:schedule-timer` *timer time &key repeat-interval absolute-p catch-up*

[Function]

Schedule *timer* to be triggered at *time*. If *absolute-p* then *time* is universal time, but non-integral values are also allowed, else *time* is measured as the number of seconds from the current time.

If *repeat-interval* is given, *timer* is automatically rescheduled upon expiry.

If *repeat-interval* is non-NIL, the Boolean *catch-up* controls whether *timer* will "catch up" by repeatedly calling its function without delay in case calls are missed because of a clock discontinuity such as a suspend and resume cycle of the computer. The default is `nil`, i.e. do not catch up.

`sb-ext:unschedule-timer` *timer*

[Function]

Cancel *timer*. Once this function returns it is guaranteed that *timer* shall not be triggered again and there are no unfinished triggers.

`sb-ext:list-all-timers`

[Function]

Return a list of all timers in the system.

15 Networking

The `sb-bsd-sockets` module provides a thinly disguised BSD socket API for SBCL. Ideas have been stolen from the BSD socket API for C and Graham Barr's `IO::Socket` classes for Perl.

Sockets are represented as CLOS objects, and the API naming conventions attempt to balance between the BSD names and good lisp style.

15.1 Sockets Overview

Most of the functions are modelled on the BSD socket API. BSD sockets are widely supported, portably (*"portable" by Unix standards, at least*) available on a variety of systems, and documented. There are some differences in approach where we have taken advantage of some of the more useful features of Common Lisp - briefly:

- Where the C API would typically return -1 and set `errno`, `sb-bsd-sockets` signals an error. All the errors are subclasses of `sb-bsd-sockets:socket-condition` and generally correspond one for one with possible `errno` values.
- We use multiple return values in many places where the C API would use pass-by-reference values.
- We can often avoid supplying an explicit *length* argument to functions because we already know how long the argument is.
- IP addresses and ports are represented in slightly friendlier fashion than "network-endian integers".

15.2 General Sockets

`sb-bsd-sockets:socket` [Class]

Class precedence list: `socket`, `standard-object`, `t`

Slots:

- `protocol` — initarg: `:protocol`; reader: `sb-bsd-sockets:socket-protocol`
Protocol used by the socket. If a keyword, the symbol-name of the keyword will be passed to `get-protocol-by-name` downcased, and the returned value used as protocol. Other values are used as-is.
- `type` — initarg: `:type`; reader: `sb-bsd-sockets:socket-type`
Type of the socket: `:stream` or `:datagram`.

Common superclass of all sockets, not meant to be directly instantiated.

`sb-bsd-sockets:socket-bind socket &rest address` [Generic Function]

Bind `socket` to `address`, which may vary according to socket family. For the `inet` family, pass `address` and `port` as two arguments; for `file` address family sockets, pass the filename string. See also `bind(2)`

`sb-bsd-sockets:socket-accept socket` [Generic Function]

Perform the `accept(2)` call, returning a newly-created connected socket and the peer address as multiple values

`sb-bsd-sockets:socket-connect socket &rest address` [Generic Function]

Perform the `connect(2)` call to connect `socket` to a remote `peer`. No useful return value.

`sb-bsd-sockets:socket-peername socket` [Generic Function]

Return SOCKET's peer; depending on the address family this may return multiple values

`sb-bsd-sockets:socket-name socket` [Generic Function]

Return the address (as vector of bytes) and port that `socket` is bound to, as multiple values.

`sb-bsd-sockets:socket-receive socket buffer length &key oob peek waitall` [Generic Function]

dontwait element-type

Read `length` octets from `socket` into `buffer` (or a freshly-consed buffer if `nil`), using `recvfrom(2)`. If `length` is `nil`, the length of `buffer` is used, so at least one of these two arguments must be non-NIL. If `buffer` is supplied, it had better be of an element type one octet wide. Returns the buffer, its length, and the address of the peer that sent it, as multiple values. On datagram sockets, sets `MSG_TRUNC` so that the actual packet length is returned even if the buffer was too small.

sb-bsd-sockets:socket-send *socket buffer length &key address external-format* [Generic Function]
oob eor dontroute dontwait nosignal confirm more

Send **length** octets from **buffer** into **socket**, using `sendto(2)`. If **buffer** is a string, it will be converted to octets according to **external-format**. If **length** is `nil`, the length of the octet buffer is used. The format of **address** depends on the socket type (for example for **inet** domain sockets it would be a list of an ip address and a port). If no socket address is provided, `send(2)` will be called instead. Returns the number of octets written.

sb-bsd-sockets:socket-listen *socket backlog* [Generic Function]

Mark **socket** as willing to accept incoming connections. The integer **backlog** defines the maximum length that the queue of pending connections may grow to before new connection attempts are refused. See also `listen(2)`.

sb-bsd-sockets:socket-open-p *socket* [Generic Function]

Return true if **socket** is open; otherwise, return false.

sb-bsd-sockets:socket-close *socket &key abort* [Generic Function]

Close **socket**, unless it was already closed.

If **socket-make-stream** has been called, calls `close` using **abort** on that stream. Otherwise closes the socket file descriptor using `close(2)`.

sb-bsd-sockets:socket-shutdown *socket &key direction* [Generic Function]

Indicate that no communication in **direction** will be performed on **socket**.

direction has to be one of `:input`, `:output` or `:io`.

After a shutdown, no input and/or output of the indicated **direction** can be performed on **socket**.

sb-bsd-sockets:socket-make-stream *socket &key input output element-type* [Generic Function]
external-format buffering timeout auto-close serve-events

Find or create a **stream** that can be used for **io** on **socket** (which must be connected). Specify whether the stream is for **input**, **output**, or both (it is an error to specify neither).

element-type and **external-format** are as per `open`.

timeout specifies a read timeout for the stream.

sb-bsd-sockets:socket-make-stream (*socket socket*) &key *input output (element-type* [Method]
(quote character)) (buffering full) (external-format default) timeout auto-close
serve-events

Default method for **socket** objects.

element-type defaults to `character`, to construct a bivalent stream, capable of both binary and character **io** use `:default`.

Acceptable values for **buffering** are `:full`, `:line` and `:none`, default is `:full`, ie. output is buffered till it is explicitly flushed using `close` or `finish-output`. (**force-output** forces some output to be flushed: to ensure all buffered output is flushed use `finish-output`.)

Streams have no **timeout** by default. If one is provided, it is the number of seconds the system will at most wait for input to appear on the socket stream when trying to read from it.

If **auto-close** is true, the underlying **os** socket is automatically closed after the stream and the socket have been garbage collected. Default is false.

If **serve-events** is true, blocking **io** on the socket will dispatch to the recursive event loop. Default is false.

The stream for **socket** will be cached, and a second invocation of this method will return the same stream. This may lead to oddities if this function is invoked with inconsistent arguments (e.g., one might request an input stream and get an output stream in response).

sb-bsd-sockets:socket-error *where &optional errno* [Function]

Signal an appropriate error for syscall **where** and **errno**.

where should be a string naming the failed function.

When supplied, **errno** should be the **unix** error number associated to the failed call. The default behavior is to use the current value of the **errno** variable.

sb-bsd-sockets:non-blocking-mode *socket* [Generic Function]

Is **socket** in non-blocking mode?

15.3 Socket Options

A subset of socket options are supported, using a fairly general framework which should make it simple to add more as required - see `SYS:CONTRIB;SB-BSD-SOCKETS:SOCKOPT.LISP` for details. The name mapping from C is fairly straightforward: `SO_RCVLOWAT` becomes `socket-receive-low-water` and (`setf socket-receive-low-water`).

- `sb-bsd-sockets:socket-reuse-address socket` [Function]
Return the value of the `so-reuseaddr` socket option for `socket`. This can also be updated with `setf`.
- `sb-bsd-sockets:socket-keep-alive socket` [Function]
Return the value of the `so-keepalive` socket option for `socket`. This can also be updated with `setf`.
- `sb-bsd-sockets:socket-oob-inline socket` [Function]
Return the value of the `so-oobinline` socket option for `socket`. This can also be updated with `setf`.
- `sb-bsd-sockets:socket-bsd-compatible socket` [Function]
Return the value of the `so-bsdcompat` socket option for `socket`. This can also be updated with `setf`. Available only on Linux.
- `sb-bsd-sockets:socket-pass-credentials socket` [Function]
Return the value of the `so-passcred` socket option for `socket`. This can also be updated with `setf`. Available only on Linux.
- `sb-bsd-sockets:socket-debug socket` [Function]
Return the value of the `so-debug` socket option for `socket`. This can also be updated with `setf`.
- `sb-bsd-sockets:socket-dont-route socket` [Function]
Return the value of the `so-dontroute` socket option for `socket`. This can also be updated with `setf`.
- `sb-bsd-sockets:socket-broadcast socket` [Function]
Return the value of the `so-broadcast` socket option for `socket`. This can also be updated with `setf`.
- `sb-bsd-sockets:socket-tcp-nodelay socket` [Function]
Return the value of the `tcp-nodelay` socket option for `socket`. This can also be updated with `setf`.

15.4 INET Domain Sockets

The TCP and UDP sockets that you know and love. Some representation issues:

- IPv4 Internet addresses are represented by vectors of (`unsigned-byte 8`) - viz. `#(127 0 0 1)`. Ports are just integers: 6010. No conversion between network- and host-order data is needed from the user of this package.
- IPv6 Internet addresses are represented by vectors of 16 (`unsigned-byte 8`) - viz. `#(0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1)`. Ports are just integers. As for IPv4 addresses, no conversion between network- and host-order data is needed from the user of this package.
- Socket addresses are represented by the two values for address and port, so for example, (`socket-connect socket #(192 168 1 1) 80`) for IPv4 and (`socket-connect socket #(0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1) 80`) for IPv6.

- `sb-bsd-sockets:inet-socket` [Class]
Class precedence list: `inet-socket`, `socket`, `standard-object`, `t`
Class representing `tcp` and `udp` over IPv4 sockets.
Examples:

```
(make-instance 'sb-bsd-sockets:inet-socket :type :stream :protocol :tcp)
```

```
(make-instance 'sb-bsd-sockets:inet-socket :type :datagram :protocol :udp)
```

- `sb-bsd-sockets:inet6-socket` [Class]
Class precedence list: `inet6-socket`, `socket`, `standard-object`, `t`
Class representing `tcp` and `udp` over IPv6 sockets.
Examples:

```
(make-instance 'sb-bsd-sockets:inet6-socket :type :stream :protocol :tcp)
```

```
(make-instance 'sb-bsd-sockets:inet6-socket :type :datagram :protocol :udp)
```


sb-bsd-sockets:make-inet-address *dotted-quads* [Function]
 Return a vector of octets given a string **dotted-quads** in the format "127.0.0.1". Signals an error if the string is malformed.

sb-bsd-sockets:make-inet6-address *colon-separated-integers* [Function]
 Return a vector of octets given a string representation of an IPv6 address **colon-separated-integers**. Signal an error if the string is malformed.

sb-bsd-sockets:get-protocol-by-name *name* [Function]
 Given a protocol name, return the protocol number, the protocol name, and a list of protocol aliases

15.5 Local (Unix) Domain Sockets

Local domain (AF_LOCAL) sockets are also known as Unix-domain sockets, but were renamed by POSIX presumably on the basis that they may be available on other systems too.

A local socket address is a string, which is used to create a node in the local filesystem. This means of course that they cannot be used across a network.

sb-bsd-sockets:local-socket [Class]
 Class precedence list: **local-socket**, **socket**, **standard-object**, **t**
 Class representing local domain (AF_LOCAL) sockets, also known as unix-domain sockets.

A local abstract socket address is also a string the scope of which is the local machine. However, in contrast to a local socket address, there is no corresponding filesystem node.

sb-bsd-sockets:local-abstract-socket [Class]
 Class precedence list: **local-abstract-socket**, **local-socket**, **socket**, **standard-object**, **t**
 Class representing local domain (AF_LOCAL) sockets with addresses in the abstract namespace.

15.6 Name Service

Presently name service is implemented by calling out to the **getaddrinfo(3)** and **gethostinfo(3)**, or to **gethostbyname(3)** **gethostbyaddr(3)** on platforms where the preferred functions are not available. The exact details of the name resolving process (for example the choice of whether DNS or a hosts file is used for lookup) are platform dependent.

sb-bsd-sockets:host-ent [Class]
 Class precedence list: **host-ent**, **standard-object**, **t**
 Slots:

- **name** — initarg: **:name**; reader: **sb-bsd-sockets:host-ent-name**
 The name of the host
- **addresses** — initarg: **:addresses**; reader: **sb-bsd-sockets:host-ent-addresses**
 A list of addresses for this host.

This class represents the results of an address lookup.

sb-bsd-sockets:get-host-by-name *node* [Function]
 Returns a **host-ent** instance for **node** or signals a **name-service-error**.
 Another **host-ent** instance containing zero, one or more IPv6 addresses may be returned as a second return value.
node may also be an **ip** address in dotted quad notation or some other weird stuff – see **getaddrinfo(3)** for the details.

sb-bsd-sockets:get-host-by-address *address* [Function]
 Returns a **host-ent** instance for **address**, which should be a vector of (integer 0 255) with 4 elements in case of an IPv4 address and 16 elements in case of an IPv6 address, or signals a **name-service-error**. See **gethostbyaddr(3)** for details.

sb-bsd-sockets:host-ent-address *host-ent* [Generic Function]
 Return some valid address for **host-ent**.

16 Profiling

SBCL includes both a deterministic profiler, that can collect statistics on individual functions, and a more “modern” statistical profiler.

Inlined functions do not appear in the results reported by either.

16.1 Deterministic Profiler

The package `sb-profile` provides a classic, per-function-call profiler.

note: When profiling code executed by multiple threads in parallel, the consing attributed to each function is inaccurate.

`sb-profile:profile` *&rest names* [Macro]
`profile` Name*

If no names are supplied, return the list of profiled functions.

If names are supplied, wrap profiling code around the named functions. As in `trace`, the names are not evaluated. A symbol names a function. A string names all the functions named by symbols in the named package. If a function is already profiled, then `unprofile` and `reprofile` (useful to notice function redefinition.) If a name is undefined, then we give a warning and ignore it. See also `unprofile`, `report` and `reset`.

`sb-profile:unprofile` *&rest names* [Macro]

Unwrap any profiling code around the named functions, or if no names are given, `unprofile` all profiled functions. A symbol names a function. A string names all the functions named by symbols in the named package. `names` defaults to the list of names of all currently profiled functions.

`sb-profile:report` *&key limit print-no-call-list* [Function]

Report results from profiling. The results are approximately adjusted for profiling overhead. The compensation may be rather inaccurate when bignums are involved in runtime calculation, as in a very-long-running Lisp process.

If `limit` is set to an integer, only the top `limit` results are reported. If `print-no-call-list` is `t` (the default) then a list of uncalled profiled functions are listed.

`sb-profile:reset` [Function]

Reset the counters for all profiled functions.

16.2 Statistical Profiler

The `sb-sprof` module, loadable by

```
(require :sb-sprof)
```

provides an alternate profiler which works by taking samples of the program execution at regular intervals, instead of instrumenting functions like `sb-profile:profile` does. You might find `sb-sprof` more useful than the deterministic profiler when profiling functions in the `common-lisp-package`, SBCL internals, or code where the instrumenting overhead is excessive.

Additionally `sb-sprof` includes a limited deterministic profiler which can be used for reporting the amounts of calls to some functions during

16.2.1 Example Usage

```
(in-package :cl-user)

(require :sb-sprof)

(declare (optimize speed))

(defun cpu-test-inner (a i)
  (logxor a
    (* i 5)
    (+ a i)))
```

```

(defun cpu-test (n)
  (let ((a 0))
    (dotimes (i (expt 2 n) a)
      (setf a (cpu-test-inner a i)))))

;;; CPU profiling

;;; Take up to 1000 samples of running (CPU-TEST 26), and give a flat
;;; table report at the end. Profiling will end once the body has been
;;; evaluated once, whether or not 1000 samples have been taken.
(sb-sprof:with-profiling (:max-samples 1000
                          :report :flat
                          :loop nil)

  (cpu-test 26))

;;; Record call counts for functions defined on symbols in the CL-USER
;;; package.
(sb-sprof:profile-call-counts "CL-USER")

;;; Take 1000 samples of running (CPU-TEST 24), and give a flat
;;; table report at the end. The body will be re-evaluated in a loop
;;; until 1000 samples have been taken. A sample count will be printed
;;; after each iteration.
(sb-sprof:with-profiling (:max-samples 1000
                          :report :flat
                          :loop t
                          :show-progress t)

  (cpu-test 24))

;;; Allocation profiling

(defun foo (&rest args)
  (mapcar (lambda (x) (float x 1d0)) args))

(defun bar (n)
  (declare (fixnum n))
  (apply #'foo (loop repeat n collect n)))

(sb-sprof:with-profiling (:max-samples 10000
                          :mode :alloc
                          :report :flat)

  (bar 1000))

```

16.2.2 Output

The flat report format will show a table of all functions that the profiler encountered on the call stack during sampling, ordered by the number of samples taken while executing that function.

Nr	Self		Total		Cumul		Calls	Function
	Count	%	Count	%	Count	%		
1	69	24.4	97	34.3	69	24.4	67108864	CPU-TEST-INNER
2	64	22.6	64	22.6	133	47.0	-	SB-VM::GENERIC-+
3	39	13.8	256	90.5	172	60.8	1	CPU-TEST
4	31	11.0	31	11.0	203	71.7	-	SB-KERNEL:TWO-ARG-XOR

For each function, the table will show three absolute and relative sample counts. The Self column shows samples taken while directly executing that function. The Total column shows samples taken while executing

that function or functions called from it (sampled to a platform-specific depth). The Cumul column shows the sum of all Self columns up to and including that line in the table.

Additionally the Calls column will record the amount of calls that were made to the function during the profiling run. This value will only be reported for functions that have been explicitly marked for call counting with `profile-call-counts`.

The profiler also hooks into the disassembler such that instructions which have been sampled are annotated with their relative frequency of sampling. This information is not stored across different sampling runs.

```

;      6CF:      702E      JO L4      ; 6/242 samples
;      6D1:      D1E3      SHL EBX, 1
;      6D3:      702A      JO L4
;      6D5: L2:  F6C303     TEST BL, 3      ; 2/242 samples
;      6D8:      756D      JNE L8
;      6DA:      8BC3      MOV EAX, EBX      ; 5/242 samples
;      6DC: L3:  83F900     CMP ECX, 0      ; 4/242 samples

```

16.2.3 Platform support

Allocation profiling is only supported on SBCL builds that use the generational garbage collector. Tracking of call stacks at a depth of more than two levels is only supported on x86 and x86-64.

16.2.4 Macros

`sb-sprof:with-profiling` (*&key sample-interval alloc-interval max-samples reset* [Macro]
mode loop max-depth show-progress threads report) *&body body*

Evaluate *body* with statistical profiling turned on. If *loop* is true, loop around the *body* until a sufficient number of samples has been collected. Returns the values from the last evaluation of *body*.

The following keyword args are recognized:

`:sample-interval` *<n>*

Take a sample every *<n>* seconds. Default is `*sample-interval*`.

`:mode` *<mode>*

If `:cpu`, run the profiler in cpu profiling mode. If `:alloc`, run the profiler in allocation profiling mode. If `:time`, run the profiler in wallclock profiling mode.

`:max-samples` *<max>*

If `:loop` is nil (the default), collect no more than *<max>* samples. If `:loop` is t, repeat evaluating *body* until *<max>* samples are taken. Default is `*max-samples*`.

`:report` *<type>*

If specified, call `report` with `:type` *<type>* at the end.

`:reset` *<bool>*

If true, call `reset` at the beginning.

`:threads` *<list-form>*

Form that evaluates to the list threads to profile, or `:all` to indicate that all threads should be profiled. Defaults to all threads.

`:threads` has no effect on call-counting at the moment.

On some platforms (eg. Darwin) the signals used by the profiler are not properly delivered to threads in proportion to their cpu usage when doing `:cpu` profiling. If you see empty call graphs, or are obviously missing several samples from certain threads, you may be falling afoul of this. In this case using `:mode :time` is likely to work better.

`:loop` *<bool>*

If false (the default), evaluate *body* only once. If true repeatedly evaluate *body*.

`sb-sprof:with-sampling` (*&optional on*) *&body body* [Macro]

Evaluate *body* with statistical sampling turned on or off in the current thread.

16.2.5 Functions

sb-sprof:map-traces *function samples*

[Function]

Call **function** on each trace in **samples**The signature of **function** must be compatible with (thread trace).**function** is called once for each trace where **thread** is the **sb-thread:tread** instance which was sampled to produce **trace**, and **trace** is an opaque object to be passed to **map-trace-pc-locs**.**experimental**: Interface subject to change.**sb-sprof:sample-pc** *info pc-or-offset*

[Function]

Extract and return program counter from **info** and **pc-or-offset**.Can be applied to the arguments passed by **map-trace-pc-locs** and **map-all-pc-locs**.**experimental**: Interface subject to change.**sb-sprof:report** *&key type max min-percent call-graph sort-by sort-order stream show-progress*

[Function]

Report statistical profiling results. The following keyword args are recognized:

:type *<type>*Specifies the type of report to generate. If **:flat**, show flat report, if **:graph** show a call graph and a flat report. If nil, don't print out a report.**:stream** *<stream>*Specify a stream to print the report on. Default is ***standard-output***.**:max** *<max>*Don't show more than *<max>* entries in the flat report.**:min-percent** *<min-percent>*Don't show functions taking less than *<min-percent>* of the total time in the flat report.**:sort-by** *<column>*If **:samples**, sort flat report by number of samples taken. If **:cumulative-samples**, sort flat report by cumulative number of samples taken (shows how much time each function spent on stack.) Default is ***report-sort-by***.**:sort-order** *<order>*If **:descending**, sort flat report in descending order. If **:ascending**, sort flat report in ascending order. Default is ***report-sort-order***.**:show-progress** *<bool>*

If true, print progress messages while generating the call graph.

:call-graph *<graph>*Print a report from *<graph>* instead of the latest profiling results.Value of this function is a **call-graph** object representing the resulting call-graph, or **nil** if there are no samples (eg. right after calling **reset**.)

Profiling is stopped before the call graph is generated.

sb-sprof:reset

[Function]

Reset the profiler.

sb-sprof:start-profiling *&key max-samples mode sample-interval alloc-interval max-depth threads*

[Function]

Start profiling statistically in the current thread if not already profiling. The following keyword args are recognized:

:sample-interval *<n>*Take a sample every *<n>* seconds. Default is ***sample-interval***.**:mode** *<mode>*If **:cpu**, run the profiler in **cpu** profiling mode. If **:alloc**, run the profiler in allocation profiling mode. If **:time**, run the profiler in wallclock profiling mode.

:max-samples <*max*>

Maximum number of stack traces to collect. Default is ***max-samples***.

:threads <*list*>

List threads to profile, or **:all** to indicate that all threads should be profiled. Defaults to **:all**.

:threads has no effect on call-counting at the moment.

On some platforms (eg. Darwin) the signals used by the profiler are not properly delivered to threads in proportion to their **cpu** usage when doing **:cpu** profiling. If you see empty call graphs, or are obviously missing several samples from certain threads, you may be falling afoul of this.

sb-sprof:stop-profiling

[Function]

Stop profiling if profiling.

sb-sprof:profile-call-counts &*rest names*

[Function]

Mark the functions named by **names** as being subject to call counting during statistical profiling. If a string is used as a name, it will be interpreted as a package name. In this case call counting will be done for all functions with names like **x** or (**setf x**), where **x** is a symbol with the package as its home package.

sb-sprof:unprofile-call-counts

[Function]

Clear all call counting information. Call counting will be done for no functions during statistical profiling.

16.2.6 Variables

sb-sprof:*max-samples*

[Variable]

Default maximum number of stack traces collected.

sb-sprof:*sample-interval*

[Variable]

Default number of seconds between samples.

16.2.7 Credits

sb-sprof is an SBCL port, with enhancements, of Gerd Moellmann's statistical profiler for CMUCL.

17 Contributed Modules

SBCL comes with a number of modules that are not part of the core system. These are loaded via `(require :modulename)` (see Section 7.9 [Customization Hooks for Users], page 59). This section contains documentation (or pointers to documentation) for some of the contributed modules.

17.1 sb-aclrepl

The **sb-aclrepl** module offers an Allegro CL-style Read-Eval-Print Loop for SBCL, with integrated inspector. Adding a debugger interface is planned.

17.1.1 Usage

To start **sb-aclrepl** as your read-eval-print loop, put the form

```
(require 'sb-aclrepl)
```

in your `~/.sbclrc` initialization file.

17.1.2 Customization

The following customization variables are available:

<code>sb-aclrepl:*command-char*</code>	[Variable]
Prefix character for a top-level command	
<code>sb-aclrepl:*prompt*</code>	[Variable]
The current prompt string or formatter function.	
<code>sb-aclrepl:*exit-on-eof*</code>	[Variable]
If <code>t</code> , then exit when the <code>eof</code> character is entered.	
<code>sb-aclrepl:*use-short-package-name*</code>	[Variable]
when <code>t</code> , use the shortest package nickname in a prompt	
<code>sb-aclrepl:*max-history*</code>	[Variable]
Maximum number of history commands to remember	

17.1.3 Example Initialization

Here's a longer example of a `~/.sbclrc` file that shows off some of the features of **sb-aclrepl**:

```
(ignore-errors (require 'sb-aclrepl))

(when (find-package 'sb-aclrepl)
  (push :aclrepl cl:*features*))
#+aclrepl
(progn
  (setq sb-aclrepl:*max-history* 100)
  (setf (sb-aclrepl:alias "asdc")
        #'(lambda (sys) (asdf:operate 'asdf:compile-op sys)))
  (sb-aclrepl:alias "l" (sys) (asdf:operate 'asdf:load-op sys))
  (sb-aclrepl:alias "t" (sys) (asdf:operate 'asdf:test-op sys))
  ;; The 1 below means that two characters ("up") are required
  (sb-aclrepl:alias ("up" 1 "Use package") (package) (use-package package))
  ;; The 0 below means only the first letter ("r") is required,
  ;; such as ":r base64"
  (sb-aclrepl:alias ("require" 0 "Require module") (sys) (require sys))
  (setf cl:*features* (delete :aclrepl cl:*features*)))
```

Questions, comments, or bug reports should be sent to Kevin Rosenberg (kevin@rosenberg.net).

17.1.4 Credits

Allegro CL is a registered trademark of Franz Inc.

17.2 sb-concurrency

Additional data structures, synchronization primitives and tools for concurrent programming. Similiar to Java's `java.util.concurrent` package.

17.2.1 Queue

`sb-concurrency:queue` is a lock-free, thread-safe FIFO queue datatype.

The implementation is based on *An Optimistic Approach to Lock-Free FIFO Queues* by Edya Ladan-Mozes and Nir Shavit.

Before SBCL 1.0.38, this implementation resided in its own contrib (see Section 17.7 [sb-queue], page 138) which is still provided for backwards-compatibility but which has since been deprecated.

`sb-concurrency:queue` [Structure]

Class precedence list: `queue`, `structure-object`, `t`

Lock-free thread safe `fifo` queue.

Use `enqueue` to add objects to the queue, and `dequeue` to remove them.

`sb-concurrency:dequeue queue` [Function]

Retrieves the oldest value in `queue` and returns it as the primary value, and `t` as secondary value. If the queue is empty, returns `nil` as both primary and secondary value.

`sb-concurrency:enqueue value queue` [Function]

Adds `value` to the end of `queue`. Returns `value`.

`sb-concurrency:list-queue-contents queue` [Function]

Returns the contents of `queue` as a list without removing them from the `queue`. Mainly useful for manual examination of queue state, as the list may be out of date by the time it is returned, and concurrent `dequeue` operations may in the worse case force the queue-traversal to be restarted several times.

`sb-concurrency:make-queue &key name initial-contents` [Function]

Returns a new `queue` with `name` and contents of the `initial-contents` sequence enqueued.

`sb-concurrency:queue-count queue` [Function]

Returns the number of objects in `queue`. Mainly useful for manual examination of queue state, and in `print-object` methods: inefficient as it must walk the entire queue.

`sb-concurrency:queue-empty-p queue` [Function]

Returns `t` if `queue` is empty, `nil` otherwise.

`sb-concurrency:queue-name instance` [Function]

Name of a `queue`. Can be assigned to using `setf`. Queue names can be arbitrary printable objects, and need not be unique.

`sb-concurrency:queuep object` [Function]

Returns true if argument is a `queue`, `nil` otherwise.

17.2.2 Mailbox (lock-free)

`sb-concurrency:mailbox` is a lock-free message queue where one or multiple ends can send messages to one or multiple receivers. The difference to [Section `sb-concurrency:queue`], page 122, is that the receiving end may block until a message arrives.

Built on top of the [Structure `sb-concurrency queue`], page 122, implementation.

<code>sb-concurrency:mailbox</code>	[Structure]
Class precedence list: <code>mailbox</code> , <code>structure-object</code> , <code>t</code>	
Mailbox aka message queue.	
<code>send-message</code> adds a message to the mailbox, <code>receive-message</code> waits till a message becomes available, whereas <code>receive-message-no-hang</code> is a non-blocking variant, and <code>receive-pending-messages</code> empties the entire mailbox in one go.	
Messages can be arbitrary objects	
<code>sb-concurrency:list-mailbox-messages mailbox</code>	[Function]
Returns a fresh list containing all the messages in the mailbox. Does not remove messages from the mailbox.	
<code>sb-concurrency:mailbox-count mailbox</code>	[Function]
Returns the number of messages currently in the mailbox.	
<code>sb-concurrency:mailbox-empty-p mailbox</code>	[Function]
Returns true if <code>mailbox</code> is currently empty, <code>nil</code> otherwise.	
<code>sb-concurrency:mailbox-name instance</code>	[Function]
Name of a mailbox. SETFable.	
<code>sb-concurrency:mailboxp object</code>	[Function]
Returns true if argument is a <code>mailbox</code> , <code>nil</code> otherwise.	
<code>sb-concurrency:make-mailbox &key name initial-contents</code>	[Function]
Returns a new <code>mailbox</code> with messages in <code>initial-contents</code> enqueued.	
<code>sb-concurrency:receive-message mailbox &key timeout</code>	[Function]
Removes the oldest message from <code>mailbox</code> and returns it as the primary value, and a secondary value of <code>t</code> . If <code>mailbox</code> is empty waits until a message arrives.	
If <code>timeout</code> is provided, and no message arrives within the specified interval, returns primary and secondary value of <code>nil</code> .	
<code>sb-concurrency:receive-message-no-hang mailbox</code>	[Function]
The non-blocking variant of <code>receive-message</code> . Returns two values, the message removed from <code>mailbox</code> , and a flag specifying whether a message could be received.	
<code>sb-concurrency:receive-pending-messages mailbox &optional n</code>	[Function]
Removes and returns all (or at most <code>n</code>) currently pending messages from <code>mailbox</code> , or returns <code>nil</code> if no messages are pending.	
Note: Concurrent threads may be snarfing messages during the run of this function, so even though <code>x,y</code> appear right next to each other in the result, does not necessarily mean that <code>y</code> was the message sent right after <code>x</code> .	
<code>sb-concurrency:send-message mailbox message</code>	[Function]
Adds a <code>message</code> to <code>mailbox</code> . Message can be any object.	

17.2.3 Gates

`sb-concurrency:gate` is a synchronization object suitable for when multiple threads must wait for a single event before proceeding.

`sb-concurrency:gate` [Structure]

Class precedence list: `gate`, `structure-object`, `t`

`gate` type. Gates are synchronization constructs suitable for making multiple threads wait for single event before proceeding.

Use `wait-on-gate` to wait for a gate to open, `open-gate` to open one, and `close-gate` to close an open gate. `gate-open-p` can be used to test the state of a gate without blocking.

`sb-concurrency:close-gate gate` [Function]

Closes `gate`. Returns `t` if the gate was previously open, and `nil` if the gate was already closed.

`sb-concurrency:gate-name instance` [Function]

Name of a `gate`. SETFable.

`sb-concurrency:gate-open-p gate` [Function]

Returns true if `gate` is open.

`sb-concurrency:gatep object` [Function]

Returns true if the argument is a `gate`.

`sb-concurrency:make-gate &key name open` [Function]

Makes a new gate. Gate will be initially open if `open` is true, and closed if `open` is `nil` (the default.) `name`, if provided, is the name of the gate, used when printing the gate.

`sb-concurrency:open-gate gate` [Function]

Opens `gate`. Returns `t` if the gate was previously closed, and `nil` if the gate was already open.

`sb-concurrency:wait-on-gate gate &key timeout` [Function]

Waits for `gate` to open, or `timeout` seconds to pass. Returns `t` if the gate was opened in time, and `nil` otherwise.

17.2.4 Frlocks, aka Fast Read Locks

sb-concurrency:frlock

[Structure]

Class precedence list: `frlock`, `structure-object`, `t`

FRlock, aka Fast Read Lock.

Fast Read Locks allow multiple readers and one potential writer to operate in parallel while providing for consistency for readers and mutual exclusion for writers.

Readers gain entry to protected regions without waiting, but need to retry if a writer operated inside the region while they were reading. This makes frlocks very efficient when readers are much more common than writers.

FRlocks are **not** suitable when it is not safe at all for readers and writers to operate on the same data in parallel: they provide consistency, not exclusion between readers and writers. Hence using an frlock to eg. protect an sbcl hash-table is unsafe. If multiple readers operating in parallel with a writer would be safe but inconsistent without a lock, frlocks are suitable.

The recommended interface to use is `frlock-read` and `frlock-write`, but those needing it can also use a lower-level interface.

Example:

```
;; Values returned by F00 are always consistent so that
;; the third value is the sum of the two first ones.
(let ((a 0)
      (b 0)
      (c 0)
      (lk (make-frlock)))
  (defun foo ()
    (frlock-read (lk) a b c))
  (defun bar (x y)
    (frlock-write (lk)
      (setf a x
            b y
            c (+ x y)))))
```

sb-concurrency:frlock-read (*frlock*) &body value-forms

[Macro]

Evaluates value-forms under frlock till it obtains a consistent set, and returns that as multiple values.

sb-concurrency:frlock-write (*frlock* &key wait-p timeout) &body body

[Macro]

Executes body while holding frlock for writing.

sb-concurrency:make-frlock &key name

[Function]

Returns a new frlock with name.

sb-concurrency:frlock-name instance

[Function]

Name of an frlock. SETFable.

sb-concurrency:frlock-read-begin frlock

[Function]

Start a read sequence on frlock. Returns a read-token and an epoch to be validated later.

Using frlock-read instead is recommended.

sb-concurrency:frlock-read-end frlock

[Function]

Ends a read sequence on frlock. Returns a token and an epoch. If the token and epoch are `eq1` to the read-token and epoch returned by `frlock-read-begin`, the values read under the frlock are consistent and can be used: if the values differ, the values are inconsistent and the read must be restated.

Using frlock-read instead is recommended.

Example:

```
(multiple-value-bind (t0 e0) (frlock-read-begin *fr*)
  (let ((a (get-a))
        (b (get-b)))
    (multiple-value-bind (t1 e1) (frlock-read-end *fr*)
      (if (and (eq1 t0 t1) (eq1 e0 e1))
          (list :a a :b b)
          :aborted)))))
```

`sb-concurrency:grab-frlock-write-lock` *frlock &key wait-p timeout* [Function]
Acquires `frlock` for writing, invalidating existing and future read-tokens for the duration. Returns `t` on success, and `nil` if the lock wasn't acquired due to eg. a timeout. Using `frlock-write` instead is recommended.

`sb-concurrency:release-frlock-write-lock` *frlock* [Function]
Releases `frlock` after writing, allowing valid read-tokens to be acquired again. Signals an error if the current thread doesn't hold `frlock` for writing. Using `frlock-write` instead is recommended.

17.3 sb-cover

The **sb-cover** module provides a code coverage tool for SBCL. The tool has support for expression coverage, and for some branch coverage. Coverage reports are only generated for code compiled using **compile-file** with the value of the **sb-cover:store-coverage-data** optimization quality set to 3.

As of SBCL 1.0.6 **sb-cover** is still experimental, and the interfaces documented here might change in later versions.

17.3.1 Example Usage

```
;;; Load SB-COVER
(require :sb-cover)

;;; Turn on generation of code coverage instrumentation in the compiler
(declare (optimize sb-cover:store-coverage-data))

;;; Load some code, ensuring that it's recompiled with the new optimization
;;; policy.
(asdf:oos 'asdf:load-op :cl-ppcre-test :force t)

;;; Run the test suite.
(cl-ppcre-test:test)

;;; Produce a coverage report
(sb-cover:report "/tmp/report/")

;;; Turn off instrumentation
(declare (optimize (sb-cover:store-coverage-data 0)))
```

17.3.2 Functions

sb-cover:report *directory &key form-mode if-matches external-format* [Function]

Print a code coverage report of all instrumented files into **directory**. If **directory** does not exist, it will be created. The main report will be printed to the file **cover-index.html**. The external format of the source files can be specified with the **external-format** parameter.

If the keyword argument **form-mode** has the value **:car**, the annotations in the coverage report will be placed on the CARs of any cons-forms, while if it has the value **:whole** the whole form will be annotated (the default). The former mode shows explicitly which forms were instrumented, while the latter mode is generally easier to read.

The keyword argument **if-matches** should be a designator for a function of one argument, called for the namestring of each file with code coverage info. If it returns true, the file's info is included in the report, otherwise ignored. The default value is **cl:identity**.

sb-cover:reset-coverage *&optional object* [Function]

Reset all coverage data back to the 'Not executed' state.

sb-cover:clear-coverage [Function]

Clear all files from the coverage database. The files will be re-entered into the database when the **fasl** files (produced by compiling **store-coverage-data** optimization policy set to 3) are loaded again into the image.

sb-cover:save-coverage [Function]

Returns an opaque representation of the current code coverage state. The only operation that may be done on the state is passing it to **restore-coverage**. The representation is guaranteed to be readably printable. A representation that has been printed and read back will work identically in **restore-coverage**.

sb-cover:save-coverage-in-file *pathname* [Function]

Call **save-coverage** and write the results of that operation into the file designated by **pathname**.

sb-cover:restore-coverage *coverage-state* [Function]

Restore the code coverage data back to an earlier state produced by **save-coverage**.

sb-cover:restore-coverage-from-file *pathname* [Function]
read the contents of the file designated by **pathname** and pass the result to **restore-coverage**.

17.4 sb-grovel

The `sb-grovel` module helps in generation of foreign function interfaces. It aids in extracting constants' values from the C compiler and in generating SB-ALIEN structure and union types, see Section 9.2.1 [Defining Foreign Types], page 71.

The ASDF(<http://www.clikli.net/ASDF>) component type GROVEL-CONSTANTS-FILE has its PERFORM operation defined to write out a C source file, compile it, and run it. The output from this program is Lisp, which is then itself compiled and loaded.

`sb-grovel` is used in a few contributed modules, and it is currently compatible only to SBCL. However, if you want to use it, here are a few directions.

17.4.1 Using sb-grovel in your own ASDF system

1. Create a Lisp package for the foreign constants/functions to go into.
2. Make your system depend on the 'sb-grovel system.
3. Create a grovel-constants data file - for an example, see `example-constants.lisp` in the `contrib/sb-grovel/` directory in the SBCL source distribution.
4. Add it as a component in your system. e.g.

```
(eval-when (:compile-toplevel :load-toplevel :execute)
  (require :sb-grovel))

(defpackage :example-package.system
  (:use :cl :asdf :sb-grovel :sb-alien))

(in-package :example-package.system)

(defsystem example-system
  :depends-on (sb-grovel)
  :components
  ((:module "sbcl"
    :components
    ((:file "defpackage")
     (grovel-constants-file "example-constants"
                           :package :example-package))))))
```

Make sure to specify the package you chose in step 1

5. Build stuff.

17.4.2 Contents of a grovel-constants-file

The `grovel-constants-file`, typically named `constants.lisp`, comprises lisp expressions describing the foreign things that you want to grovel for. A `constants.lisp` file contains two sections:

- a list of headers to include in the C program, for example:

```
("sys/types.h" "sys/socket.h" "sys/stat.h" "unistd.h" "sys/un.h"
 "netinet/in.h" "netinet/in_systm.h" "netinet/ip.h" "net/if.h"
 "netdb.h" "errno.h" "netinet/tcp.h" "fcntl.h" "signal.h" )
```

- A list of sb-grovel clauses describing the things you want to grovel from the C compiler, for example:

```
((:integer af-local
  #+(or sunos solaris) "AF_UNIX"
  #-(or sunos solaris) "AF_LOCAL"
  "Local to host (pipes and file-domain).")
 (:structure stat ("struct stat"
  (integer dev "dev_t" "st_dev")
  (integer atime "time_t" "st_atime")))
 (:function getpid ("getpid" int )))
```

There are two types of things that `sb-grovel` can sensibly extract from the C compiler: constant integers and structure layouts. It is also possible to define foreign functions in the `constants.lisp` file, but these

definitions don't use any information from the C program; they expand directly to `sb-alien:define-alien-routine` (see Section 9.7.2 [The `define-alien-routine` Macro], page 78) forms.

Here's how to use the grovel clauses:

- `:integer` - constant expressions in C. Used in this form:

```
(:integer lisp-variable-name "C expression" &optional doc export)
```

"C expression" will be typically be the name of a constant. But other forms are possible.
- `:enum`

```
(:enum lisp-type-name ((lisp-enumerated-name c-enumerated-name) ...)))
```

An `sb-alien:enum` type with name `lisp-type-name` will be defined. The symbols are the `lisp-enumerated-names`, and the values are grovelled from the `c-enumerated-names`.

- `:structure` - alien structure definitions look like this:

```
(:structure lisp-struct-name ("struct c_structure"
                             (type-designator lisp-element-name
                              "c_element_type" "c_element_name"
                              :distrust-length nil)
                             ; ...
                             ))
```

`type-designator` is a reference to a type whose size (and type constraints) will be grovelled for. `sb-grovel` accepts a form of type designator that doesn't quite conform to either lisp nor `sb-alien`'s type specifiers. Here's a list of type designators that `sb-grovel` currently accepts:

- `integer` - a C integral type; `sb-grovel` will infer the exact type from size information extracted from the C program. All common C integer types can be grovelled for with this type designator, but it is not possible to grovel for bit fields yet.
- `(unsigned n)` - an unsigned integer variable that is `n` bytes long. No size information from the C program will be used.
- `(signed n)` - a signed integer variable that is `n` bytes long. No size information from the C program will be used.
- `c-string` - an array of `char` in the structure. `sb-grovel` will use the array's length from the C program, unless you pass it the `:distrust-length` keyword argument with non-`nil` value (this might be required for structures such as `solaris's struct dirent`).
- `c-string-pointer` - a pointer to a C string, corresponding to the `sb-alien:c-string` type (see Section 9.2.3 [Foreign Type Specifiers], page 72).
- `(array alien-type)` - An array of the previously-declared alien type. The array's size will be determined from the output of the C program and the alien type's size.
- `(array alien-type n)` - An array of the previously-declared alien type. The array's size will be assumed as being `n`.

Note that `c-string` and `c-string-pointer` do not have the same meaning. If you declare that an element is of type `c-string`, it will be treated as if the string is a part of the structure, whereas if you declare that the element is of type `c-string-pointer`, a *pointer to a string* will be the structure member.

- `:function` - alien function definitions are similar to `define-alien-routine` definitions, because they expand to such forms when the lisp program is loaded. See Section 9.7 [Foreign Function Calls], page 78.

```
(:function lisp-function-name ("alien_function_name" alien-return-type
                              (argument alien-type)
                              (argument2 alien-type)))
```

17.4.3 Programming with `sb-grovel`'s structure types

Let us assume that you have a grovelled structure definition:

```
(:structure mystruct ("struct my_structure"
                     (integer myint "int" "st_int")
                     (c-string mystring "char[]" "st_str")))
```

What can you do with it? Here's a short interface document:

- Creating and destroying objects:
 - Function (`allocate-mystruct`) - allocates an object of type `mystruct` and returns a system area pointer to it.
 - Function (`free-mystruct var`) - frees the alien object pointed to by `var`.
 - Macro (`with-mystruct var ((member init) [...]) &body body`) - allocates an object of type `mystruct` that is valid in `body`. If `body` terminates or control unwinds out of `body`, the object pointed to by `var` will be deallocated.
- Accessing structure members:
 - (`mystruct-myint var`) and (`mystruct-mystring var`) return the value of the respective fields in `mystruct`.
 - (`setf (mystruct-myint var) new-val`) and (`setf (mystruct-mystring var) new-val`) sets the value of the respective structure member to the value of `new-val`. Notice that in (`setf (mystruct-mystring var) new-val`)'s case, `new-val` is a lisp string.

17.4.3.1 Traps and Pitfalls

Basically, you can treat functions and data structure definitions that `sb-grovel` spits out as if they were alien routines and types. This has a few implications that might not be immediately obvious (especially if you have programmed in a previous version of `sb-grovel` that didn't use alien types):

- You must take care of `grovel`-allocated structures yourself. They are alien types, so the garbage collector will not collect them when you drop the last reference.
- If you use the `with-mystruct` macro, be sure that no references to the variable thus allocated leaks out. It will be deallocated when the block exits.

17.5 sb-md5

The **sb-md5** module implements the RFC1321 MD5 Message Digest Algorithm. [FIXME cite]

sb-md5:md5sum-file *pathname* [Function]

Calculate the MD5 message-digest of the file specified by ‘pathname’.

sb-md5:md5sum-sequence *sequence &key start end* [Function]

Calculate the MD5 message-digest of data in ‘sequence’, which should be a 1d simple-array with element type (unsigned-byte 8). On **cmu cl** and **sbcl** non-simple and non-1d arrays with this element-type are also supported.

sb-md5:md5sum-stream *stream* [Function]

Calculate an MD5 message-digest of the contents of ‘stream’. Its element-type has to be (unsigned-byte 8). Use on character streams is **deprecated**, as this will not work correctly on implementations with ‘char-code-limit’ > 256 and ignores character coding issues.

sb-md5:md5sum-string *string &key external-format start end* [Function]

Calculate the MD5 message-digest of the binary representation of ‘string’ (as octets) in the external format specified by ‘external-format’. The boundaries ‘start’ and ‘end’ refer to character positions in the string, not to octets in the resulting binary representation. The permissible external format specifiers are determined by the underlying implementation.

17.5.1 Credits

The implementation for CMUCL was largely done by Pierre Mai, with help from members of the **cmucl-help** mailing list. Since CMUCL and SBCL are similar in many respects, it was not too difficult to extend the low-level implementation optimizations for CMUCL to SBCL. Following this, SBCL’s compiler was extended to implement efficient compilation of modular arithmetic (see Section 6.3 [Modular arithmetic], page 39), which enabled the implementation to be expressed in portable arithmetical terms, apart from the use of **rotate-byte** for bitwise rotation.

17.6 sb-posix

Sb-posix is the supported interface for calling out to the operating system.¹

The scope of this interface is “operating system calls on a typical Unixlike platform”. This is section 2 of the Unix manual, plus section 3 calls that are (a) typically found in `libc`, but (b) not part of the C standard. For example, we intend to provide support for `opendir()` and `readdir()`, but not for `printf()`. That said, if your favourite system call is not included yet, you are encouraged to submit a patch to the SBCL mailing list.

Some facilities are omitted where they offer absolutely no additional use over some portable function, or would be actively dangerous to the consistency of Lisp. Not all functions are available on all platforms.

17.6.1 Lisp names for C names

All symbols are in the `SB-POSIX` package. This package contains a Lisp function for each supported Unix system call or function, a variable or constant for each supported Unix constant, an object type for each supported Unix structure type, and a slot name for each supported Unix structure member. A symbol name is derived from the C binding’s name, by (a) uppercasing, then (b) removing leading underscores (`#_`) then replacing remaining underscore characters with the hyphen (`#\-`). The requirement to uppercase is so that in a standard upcasing reader the user may write `sb-posix:creat` instead of `sb-posix:|creat|` as would otherwise be required.

No other changes to “Lispify” symbol names are made, so `creat()` becomes `CREATE`, not `CREATEE`.

The user is encouraged not to `(USE-PACKAGE :SB-POSIX)` but instead to use the `SB-POSIX:` prefix on all references, as some of the symbols contained in the `SB-POSIX` package have the same name as CL symbols (`OPEN`, `CLOSE`, `SIGNAL` etc).

17.6.2 Types

Generally, marshalling between Lisp and C data types is done using SBCL’s FFI. See Chapter 9 [Foreign Function Interface], page 71.

Some functions accept objects such as filenames or file descriptors. In the C binding to POSIX these are represented as strings and small integers respectively. For the Lisp programmer’s convenience we introduce designators such that CL pathnames or open streams can be passed to these functions. For example, `rename` accepts both pathnames and strings as its arguments.

17.6.2.1 File-descriptors

`sb-posix:file-descriptor` [Type]

A `fixnum` designating a native file descriptor.

`sb-sys:make-fd-stream` can be used to construct a `file-stream` associated with a native file descriptor. Note that mixing I/O operations on a `file-stream` with operations directly on its descriptor may produce unexpected results if the stream is buffered.

`sb-posix:file-descriptor-designator` [Type]

Designator for a `file-descriptor`: either a `fixnum` designating itself, or a `file-stream` designating the underlying file-descriptor.

`sb-posix:file-descriptor` *file-descriptor* [Function]

Converts `file-descriptor-designator` into a `file-descriptor`.

17.6.2.2 Filenames

`sb-posix:filename` [Type]

A `string` designating a filename in native namestring syntax.

Note that native namestring syntax is distinct from Lisp namestring syntax:

`(pathname "/foo*/bar")`

is a wild pathname with a pattern-matching directory component. `sb-ext:parse-native-namestring` may be used to construct Lisp pathnames that denote `posix` filenames as understood by system calls, and `sb-ext:native-namestring` can be used to coerce them into strings in the native namestring syntax.

¹ The functionality contained in the package `SB-UNIX` is for SBCL internal use only; its contents are likely to change from version to version.

Note also that `posix` filename syntax does not distinguish the names of files from the names of directories: in order to parse the name of a directory in `posix` filename syntax into a pathname `my-defaults` for which

```
(merge-pathnames (make-pathname :name "FOO" :case :common)
                  my-defaults)
```

returns a pathname that denotes a file in the directory, supply a true `:as-directory` argument to `sb-ext:parse-native-namestring`. Likewise, to supply the name of a directory to a `posix` function in non-directory syntax, supply a true `:as-file` argument to `sb-ext:native-namestring`.

`sb-posix:filename-designator` [Type]
Designator for a `filename`: a `string` designating itself, or a designator for a `pathname` designating the corresponding native namestring.

`sb-posix:filename filename` [Function]
Converts `filename-designator` into a `filename`.

17.6.3 Function Parameters

The calling convention is modelled after that of CMUCL's UNIX package: in particular, it's like the C interface except that:

- Length arguments are omitted or optional where the sensible value is obvious. For example, `read` would be defined this way:


```
(read fd buffer &optional (length (length buffer))) => bytes-read
```
- Where C simulates “out” parameters using pointers (for instance, in `pipe()` or `socketpair()`) these may be optional or omitted in the Lisp interface: if not provided, appropriate objects will be allocated and returned (using multiple return values if necessary).
- Some functions accept objects such as filenames or file descriptors. Wherever these are specified as such in the C bindings, the Lisp interface accepts designators for them as specified in the ‘Types’ section above.
- A few functions have been included in `sb-posix` that do not correspond exactly with their C counterparts. These are described in See Section 17.6.6 [Functions with idiosyncratic bindings], page 137.

17.6.4 Function Return Values

The return value is usually the same as for the C binding, except in error cases: where the C function is defined as returning some sentinel value and setting `errno` on error, we instead signal an error of type `SYSCALL-ERROR`. The actual error value (`errno`) is stored in this condition and can be accessed with `SYSCALL-ERRNO`.

We do not automatically translate the returned value into “Lispy” objects – for example, `SB-POSIX:OPEN` returns a small integer, not a stream. Exception: boolean-returning functions (or, more commonly, macros) do not return a C integer, but instead a Lisp boolean.

17.6.5 Lisp objects and C structures

`Sb-posix` provides various Lisp object types to stand in for C structures in the POSIX library. Lisp bindings to C functions that accept, manipulate, or return C structures accept, manipulate, or return instances of these Lisp types instead of instances of alien types.

The names of the Lisp types are chosen according to the general rules described above. For example Lisp objects of type `STAT` stand in for C structures of type `struct stat`.

Accessors are provided for each standard field in the structure. These are named `structure-name-field-name` where the two components are chosen according to the general name conversion rules, with the exception that in cases where all fields in a given structure have a common prefix, that prefix is omitted. For example, `stat.st_dev` in C becomes `STAT-DEV` in Lisp.

Because `sb-posix` might not support all semi-standard or implementation-dependent members of all structure types on your system (patches welcome), here is an enumeration of all supported Lisp objects corresponding to supported POSIX structures, and the supported slots for those structures.

- `flock`

`sb-posix:flock` [Class]
Class precedence list: `flock`, `standard-object`, `t`

Slots:

- **type** — initarg: `:type`; reader: `sb-posix:flock-type`; writer: `(setf sb-posix:flock-type)`
Type of lock; `F_RDLCK`, `F_WRLCK`, `F_UNLCK`.
- **whence** — initarg: `:whence`; reader: `sb-posix:flock-whence`; writer: `(setf sb-posix:flock-whence)`
Flag for starting offset.
- **start** — initarg: `:start`; reader: `sb-posix:flock-start`; writer: `(setf sb-posix:flock-start)`
Relative offset in bytes.
- **len** — initarg: `:len`; reader: `sb-posix:flock-len`; writer: `(setf sb-posix:flock-len)`
Size; if 0 then until eof.
- **pid** — reader: `sb-posix:flock-pid`
Process id of the process holding the lock; returned with `F_GETLK`.

Class representing locks used in `fcntl(2)`.

- **passwd**

`sb-posix:passwd` [Class]

Class precedence list: `passwd`, `standard-object`, `t`

Slots:

- **name** — initarg: `:name`; reader: `sb-posix:passwd-name`; writer: `(setf sb-posix:passwd-name)`
User's login name.
- **passwd** — initarg: `:passwd`; reader: `sb-posix:passwd-passwd`; writer: `(setf sb-posix:passwd-passwd)`
The account's encrypted password.
- **uid** — initarg: `:uid`; reader: `sb-posix:passwd-uid`; writer: `(setf sb-posix:passwd-uid)`
Numerical user id.
- **gid** — initarg: `:gid`; reader: `sb-posix:passwd-gid`; writer: `(setf sb-posix:passwd-gid)`
Numerical group id.
- **gecos** — initarg: `:gecos`; reader: `sb-posix:passwd-gecos`; writer: `(setf sb-posix:passwd-gecos)`
User's name or comment field.
- **dir** — initarg: `:dir`; reader: `sb-posix:passwd-dir`; writer: `(setf sb-posix:passwd-dir)`
Initial working directory.
- **shell** — initarg: `:shell`; reader: `sb-posix:passwd-shell`; writer: `(setf sb-posix:passwd-shell)`
Program to use as shell.

Instances of this class represent entries in the system's user database.

- **stat**

`sb-posix:stat` [Class]

Class precedence list: `stat`, `standard-object`, `t`

Slots:

- **mode** — initarg: `:mode`; reader: `sb-posix:stat-mode`
Mode of file.
- **ino** — initarg: `:ino`; reader: `sb-posix:stat-ino`
File serial number.
- **dev** — initarg: `:dev`; reader: `sb-posix:stat-dev`
Device id of device containing file.

- **nlink** — initarg: `:nlink`; reader: `sb-posix:stat-nlink`
Number of hard links to the file.
- **uid** — initarg: `:uid`; reader: `sb-posix:stat-uid`
User id of file.
- **gid** — initarg: `:gid`; reader: `sb-posix:stat-gid`
Group id of file.
- **size** — initarg: `:size`; reader: `sb-posix:stat-size`
For regular files, the file size in bytes. For symbolic links, the length in bytes of the filename contained in the symbolic link.
- **rdev** — initarg: `:rdev`; reader: `sb-posix:stat-rdev`
For devices the device number.
- **atime** — initarg: `:atime`; reader: `sb-posix:stat-atime`
Time of last access.
- **mtime** — initarg: `:mtime`; reader: `sb-posix:stat-mtime`
Time of last data modification.
- **ctime** — initarg: `:ctime`; reader: `sb-posix:stat-ctime`
Time of last status change.

Instances of this class represent `posix` file metadata.

- **termios**

`sb-posix:termios` [Class]

Class precedence list: `termios`, `standard-object`, `t`

Slots:

- **iflag** — initarg: `:iflag`; reader: `sb-posix:termios-iflag`; writer: `(setf sb-posix:termios-iflag)`
Input modes.
- **oflag** — initarg: `:oflag`; reader: `sb-posix:termios-oflag`; writer: `(setf sb-posix:termios-oflag)`
Output modes.
- **cflag** — initarg: `:cflag`; reader: `sb-posix:termios-cflag`; writer: `(setf sb-posix:termios-cflag)`
Control modes.
- **lflag** — initarg: `:lflag`; reader: `sb-posix:termios-lflag`; writer: `(setf sb-posix:termios-lflag)`
Local modes.
- **cc** — initarg: `:cc`; reader: `sb-posix:termios-cc`; writer: `(setf sb-posix:termios-cc)`
Control characters.

Instances of this class represent I/O characteristics of the terminal.

- **timeval**

`sb-posix:timeval` [Class]

Class precedence list: `timeval`, `standard-object`, `t`

Slots:

- **sec** — initarg: `:tv-sec`; reader: `sb-posix:timeval-sec`; writer: `(setf sb-posix:timeval-sec)`
Seconds.
- **usec** — initarg: `:tv-usec`; reader: `sb-posix:timeval-usec`; writer: `(setf sb-posix:timeval-usec)`
Microseconds.

Instances of this class represent time values.

17.6.6 Functions with idiosyncratic bindings

A few functions in `sb-posix` don't correspond directly to their C counterparts.

- `getcwd`

`sb-posix:cwd`

[Function]

Returns the process's current working directory as a string.

- `readlink`

`sb-posix:readlink pathspec`

[Function]

Returns the resolved target of a symbolic link as a string.

- `syslog`

`sb-posix:syslog priority format &rest args`

[Function]

Send a message to the syslog facility, with severity level `priority`. The message will be formatted as by `cl:format` (rather than C's `printf`) with format string `format` and arguments `args`.

17.7 **sb-queue**

Since SBCL 1.0.38, the **sb-queue** module has been merged into the **sb-concurrency** module (see Section 17.2 [sb-concurrency], page 121.)

17.8 sb-rotate-byte

The **sb-rotate-byte** module offers an interface to bitwise rotation, with an efficient implementation for operations which can be performed directly using the platform's arithmetic routines. It implements the specification at <http://www.clike.net/ROTATE-BYTE>.

Bitwise rotation is a component of various cryptographic or hashing algorithms: MD5, SHA-1, etc.; often these algorithms are specified on 32-bit rings. [FIXME cite cite cite].

sb-rotate-byte:rotate-byte *count bytespec integer* [Function]
Rotates a field of bits within **integer**; specifically, returns an integer that contains the bits of **integer** rotated **count** times leftwards within the byte specified by **bytespec**, and elsewhere contains the bits of **integer**.

18 Deprecation

In order to support evolution of interfaces in SBCL as well as in user code, SBCL allows declaring functions, variables and types as deprecated. Users of deprecated things are notified by means of warnings while the deprecated thing in question is still available.

This chapter documents the interfaces for being notified when using deprecated thing and declaring things as deprecated, the deprecation process used for SBCL interfaces, and lists legacy interfaces in various stages of deprecation.

Deprecation in this context should not be confused with those things the ANSI Common Lisp standard calls *deprecated*: the entirety of ANSI CL is supported by SBCL, and none of those interfaces are subject to censure.

18.1 Why Deprecate?

While generally speaking we try to keep SBCL changes as backwards compatible as feasible, there are situations when existing interfaces are deprecated:

- **Broken Interfaces**

Sometimes it turns out that an interface is sufficiently misdesigned that fixing it would be worse than deprecating it and replacing it with another.

This is typically the case when fixing the interface would change its semantics in ways that could break user code subtly: in such cases we may end up considering the obvious breakage caused by deprecation to be preferable.

Another example are functions or macros whose current signature makes them hard or impossible to extend in the future: backwards compatible extensions would either make the interface intolerably hairy, or are sometimes outright impossible.

- **Internal Interfaces**

SBCL has several internal interfaces that were never meant to be used in user code – or at least never meant to be used in user code unwilling to track changes to SBCL internals.

Ideally we'd like to be free to refactor our own internals as we please, without even going through the hassle of deprecating things. Sometimes, however, it turns out that our internal interfaces have several external users who aren't using them advisedly, but due to misunderstandings regarding their status or stability.

Consider a deprecated internal interface a reminder for SBCL maintainers not to delete the thing just yet, even though it seems unused – because it has external users.

When internal interfaces are deprecated we try our best to provide supported alternatives.

- **Aesthetics & Ease of Maintenance**

Sometimes an interface isn't broken or internal, but just inconsistent somehow.

This mostly happens only with historical interfaces inherited from CMUCL which often haven't been officially supported in SBCL before, or with new extensions to SBCL that haven't been around for very long in the first place.

The alternative would be to keep the suboptimal version around forever, possibly alongside an improved version. Sometimes we may do just that, but because every line of code comes with a maintenance cost, sometimes we opt to deprecate the suboptimal version instead: SBCL doesn't have infinite developer resources.

We also believe that sometimes cleaning out legacy interfaces helps keep the whole system more comprehensible to users, and makes introspective tools such as `apropos` more useful.

18.2 The Deprecation Pipeline

SBCL uses a *deprecation pipeline* with multiple stages: as time goes by, deprecated things move from earlier stages of deprecation to later stages before finally being removed. The intention is making users aware of necessary changes early but allowing a migration to new interfaces at a reasonable pace.

Deprecation proceeds in three stages, each lasting approximately a year. In some cases it might move slower or faster, but one year per stage is what we aim at in general. During each stage warnings (and errors)

of increasing severity are signaled, which note that the interface is deprecated, and point users towards any replacements when applicable.

1. Early Deprecation

During early deprecation the interface is kept in working condition. However, when a thing in this deprecation stage is used, an `sb-ext:early-deprecation-warning`, which is a `style-warning`, is signaled at compile-time.

The internals may change at this stage: typically because the interface is re-implemented on top of its successor. While we try to keep things as backwards-compatible as feasible (taking maintenance costs into account), sometimes semantics change slightly.

For example, when the spinlock API was deprecated, spinlock objects ceased to exist, and the whole spinlock API became a synonym for the mutex API – so code using the spinlock API continued working, but silently switched to mutexes instead. However, if someone relied on

```
(typep lock 'spinlock)
```

returning NIL for a mutexes, trouble could ensue.

2. Late Deprecation

During late deprecation the interface remains as it was during early deprecation, but the compile-time warning is upgraded: when a thing in this deprecation stage is used, a `sb-ext:late-deprecation-warning`, which is a full `warning`, is signaled at compile-time.

3. Final Deprecation

During final deprecation the symbols still exist. However, when a thing in this deprecation stage is used, a `sb-ext:final-deprecation-warning`, which is a full `warning`, is signaled at compile-time and an `error` is signaled at run-time.

4. After Final Deprecation

The interface is deleted entirely.

18.3 Deprecation Conditions

`sb-ext:deprecation-condition` is the superclass of all deprecation-related warning and error conditions. All common slots and readers are defined in this condition class.

`sb-ext:deprecation-condition` [Condition]

Class precedence list: `deprecation-condition`, `condition`, `t`

Superclass for deprecation-related error and warning conditions.

`sb-ext:early-deprecation-warning` [Condition]

Class precedence list: `early-deprecation-warning`, `style-warning`, `warning`, `deprecation-condition`, `condition`, `t`

This warning is signaled when the use of a variable, function, type, etc. in `:early` deprecation is detected at compile-time. The use will work at run-time with no warning or error.

`sb-ext:late-deprecation-warning` [Condition]

Class precedence list: `late-deprecation-warning`, `warning`, `deprecation-condition`, `condition`, `t`

This warning is signaled when the use of a variable, function, type, etc. in `:late` deprecation is detected at compile-time. The use will work at run-time with no warning or error.

`sb-ext:final-deprecation-warning` [Condition]

Class precedence list: `final-deprecation-warning`, `warning`, `deprecation-condition`, `condition`, `t`

This warning is signaled when the use of a variable, function, type, etc. in `:final` deprecation is detected at compile-time. An error will be signaled at run-time.

`sb-ext:deprecation-error` [Condition]

Class precedence list: `deprecation-error`, `error`, `serious-condition`, `deprecation-condition`, `condition`, `t`

This error is signaled at run-time when an attempt is made to use a thing that is in `:final` deprecation, i.e. call a function or access a variable.

18.4 Introspecting Deprecation Information

The deprecation status of functions and variables can be inspected using the `sb-cltl2:function-information` and `sb-cltl2:variable-information` functions provided by the `sb-cltl2` contributed module.

18.5 Deprecation Declaration

The `sb-ext:deprecated` declaration can be used to declare objects in various namespaces¹ as deprecated.

`sb-ext:deprecated` [Declaration]
Syntax:

```
sb-ext:deprecated stage since {object-clause}*

stage ::= { :early | :late | :final }

since ::= { version | (software version) }

object-clause ::= (namespace name [:replacement replacement])

namespace ::= { cl:variable | cl:function | cl:type }
```

where *name* is the name of the deprecated thing, *version* and *software* are strings describing the version in which the thing has been deprecated and *replacement* is a name or a list of names designating things that should be used instead of the deprecated thing.

Currently the following namespaces are supported:

cl:function

Declare functions, compiler-macros or macros as deprecated.

note: When declaring a function to be in `:final` deprecation, there should be no actual definition of the function as the declaration emits a stub function that signals a `sb-ext:deprecation-error` at run-time when called.

cl:variable

Declare special and global variables, constants and symbol-macros as deprecated.

note: When declaring a variable to be in `:final` deprecation, there should be no actual definition of the variable as the declaration emits a symbol-macro that signals a `sb-ext:deprecation-error` at run-time when accessed.

cl:type Declare named types (i.e. defined via `deftype`), standard classes, structure classes and condition classes as deprecated.

18.6 Deprecation Examples

Marking functions as deprecated:

```
(defun foo ())
(defun bar ())
(declare (deprecated :early ("my-system" "1.2.3")
                     (function foo :replacement bar)))

;; Remember: do not define the actual function or variable in case of
;; :final deprecation:
(declare (deprecated :final ("my-system" "1.2.3")
                    (function fez :replacement whoop)))
```

Attempting to use the deprecated functions:

```
(defun baz ()
  (foo))
| STYLE-WARNING: The function CL-USER::FOO has been deprecated...
```

¹ See “namespace” entry in the glossary of the Common Lisp Hyperspec.

```

=> BAZ
(baz)
=> NIL ; no error

(defun danger ()
  (fez))
| WARNING: The function CL-USER::FEZ has been deprecated...
=> DANGER
(danger)
|- ERROR: The function CL-USER::FEZ has been deprecated...

```

18.7 Deprecated Interfaces in SBCL

This sections lists legacy interfaces in various stages of deprecation.

18.7.1 List of Deprecated Interfaces

18.7.1.1 Early Deprecation

- **SOCKINT::WIN32-***

Deprecated in favor of the corresponding prefix-less functions (e.g. `socket::bind` replaces `socketint::win32-bind`) as of 1.2.10 in March 2015. Expected to move into late deprecation in August 2015.

- **SB-UNIX:UNIX-EXIT**

Deprecated as of 1.0.56.55 in May 2012. Expected to move into late deprecation in May 2013.

When the SBCL process termination was refactored as part of changes that led to `sb-ext:quit` being deprecated, `sb-unix:unix-exit` ceased to be used internally. Since **SB-UNIX** is an internal package not intended for user code to use, and since we're slowly in the process of refactoring things to be less Unix-oriented, `sb-unix:unix-exit` was initially deleted as it was no longer used. Unfortunately it became apparent that it was used by several external users, so it was re-instated in deprecated form.

While the cost of keeping `sb-unix:unix-exit` indefinitely is trivial, the ability to refactor our internals is important, so its deprecation was taken as an opportunity to highlight that **SB-UNIX** is an internal package and **SB-POSIX** should be used by user-programs instead – or alternatively calling the foreign function directly if the desired interface doesn't for some reason exist in **SB-POSIX**.

Remedy

For code needing to work with legacy SBCLs, use e.g. `system-exit` as show above in remedies for `sb-ext:quit`. In modern SBCLs simply call either `sb-posix:exit` or `sb-ext:exit` with appropriate arguments.

- **SB-C::MERGE-TAIL-CALLS Compiler Policy**

Deprecated as of 1.0.53.74 in November 2011. Expected to move into late deprecation in November 2012.

This compiler policy was never functional: SBCL has always merged tail calls when it could, regardless of this policy setting. (It was also never officially supported, but several code-bases have historically used it.)

Remedy

Simply remove the policy declarations. They were never necessary: SBCL always merged tail-calls when possible. To disable tail merging, structure the code to avoid the tail position instead.

- **Spinlock API**

Deprecated as of 1.0.53.11 in August 2011. Expected to move into late deprecation in August 2012.

Spinlocks were an internal interface, but had a number of external users and were hence deprecated instead of being simply deleted.

Affected symbols: `sb-thread::spinlock`, `sb-thread::make-spinlock`, `sb-thread::with-spinlock`, `sb-thread::with-recursive-spinlock`, `sb-thread::get-spinlock`, `sb-thread::release-spinlock`, `sb-thread::spinlock-value`, and `sb-thread::spinlock-name`.

Remedy

Use the mutex API instead, or implement spinlocks suiting your needs on top of `sb-ext:compare-and-swap`, `sb-ext:spin-loop-hint`, etc.

- **SOCKINT::HANDLE->FD, SOCKINT::FD->HANDLE**

Internally deprecated in 2012. Declared deprecated as of 1.2.10 in March 2015. Expected to move into final deprecation in August 2015.

18.7.1.2 Late Deprecation

- **SB-THREAD:JOIN-THREAD-ERROR-THREAD and SB-THREAD:INTERRUPT-THREAD-ERROR-THREAD**

Deprecated in favor of `sb-thread:thread-error-thread` as of 1.0.29.17 in June 2009. Expected to move into final deprecation in June 2012.

Remedy

For code that needs to support legacy SBCLs, use e.g.:

```
(defun get-thread-error-thread (condition)
  #+#. (cl:if (cl:find-symbol "THREAD-ERROR-THREAD" :sb-thread)
    '(and) '(or))
  (sb-thread:thread-error-thread condition)
  #-#. (cl:if (cl:find-symbol "THREAD-ERROR-THREAD" :sb-thread)
    '(and) '(or))
  (etypecase condition
    (sb-thread:join-thread-error
     (sb-thread:join-thread-error-thread condition))
    (sb-thread:interrupt-thread-error
     (sb-thread:interrupt-thread-error-thread condition))))
```

- **SB-INTROSPECT:FUNCTION-ARGLIST**

Deprecated in favor of `sb-introspect:function-lambda-list` as of 1.0.24.5 in January 2009. Expected to move into final deprecation in January 2012.

Renamed for consistency and aesthetics. Functions have lambda-lists, not arglists.

Remedy

For code that needs to support legacy SBCLs, use e.g.:

```
(defun get-function-lambda-list (function)
  #+#. (cl:if (cl:find-symbol "FUNCTION-LAMBDA-LIST" :sb-introspect)
    '(and) '(or))
  (sb-introspect:function-lambda-list function)
  #-#. (cl:if (cl:find-symbol "FUNCTION-LAMBDA-LIST" :sb-introspect)
    '(and) '(or))
  (sb-introspect:function-arglist function))
```

- **Stack Allocation Policies**

Deprecated in favor of `sb-ext:*stack-allocate-dynamic-extent*` as of 1.0.19.7 in August 2008, and are expected to be removed in August 2012.

Affected symbols: `sb-c::stack-allocate-dynamic-extent`, `sb-c::stack-allocate-vector`, and `sb-c::stack-allocate-value-cells`.

These compiler policies were never officially supported, and turned out to be a flawed design.

Remedy

For code that needs stack-allocation in legacy SBCLs, conditionalize using:

```
#-#. (cl:if (cl:find-symbol "*STACK-ALLOCATE-DYNAMIC-EXTENT*" :sb-ext)
           '(and) '(or))
(declare (optimize sb-c::stack-allocate-dynamic-extent))
```

However, unless stack allocation is essential, we recommend simply removing these declarations. Refer to documentation on `sb-ext:*stack-allocate-dynamic*` for details on stack allocation control in modern SBCLs.

- **SB-SYS:OUTPUT-RAW-BYTES**

Deprecated as of 1.0.8.16 in June 2007. Expected to move into final deprecation in June 2012.

Internal interface with some external users. Never officially supported, deemed unnecessary in presence of `write-sequence` and bivalent streams.

Remedy

Use streams with element-type (`unsigned-byte 8`) or `:default` – the latter allowing both binary and character IO – in conjunction with `write-sequence`.

18.7.1.3 Final Deprecation

No interfaces are currently in final deprecation.

18.7.2 Historical Interfaces

The following is a partial list of interfaces present in historical versions of SBCL, which have since then been deleted.

- **SB-KERNEL:INSTANCE-LAMBDA**

Historically needed for CLOS code. Deprecated as of 0.9.3.32 in August 2005. Deleted as of 1.0.47.8 in April 2011. Plain `lambda` can be used where `sb-kernel:instance-lambda` used to be needed.

- **SB-ALIEN:DEF-ALIEN-ROUTINE, SB-ALIEN:DEF-ALIEN-VARIABLE, SB-ALIEN:DEF-ALIEN-TYPE**

Inherited from CMUCL, naming convention not consistent with preferred SBCL style. Deprecated as of 0.pre7.90 in December 2001. Deleted as of 1.0.9.17 in September 2007. Replaced by `sb-alien:define-alien-routine`, `sb-alien:define-alien-variable`, and `sb-alien:define-alien-type`.

Appendix A Concept Index

A

Actual Source 17, 18
 Arithmetic, hardware 39, 139
 Arithmetic, modular 39, 139
 Asynchronous Timeout 64
 Availability of debug variables 30

B

Block compilation, debugger implications 28
 Block, basic 32
 Block, start location 32

C

Character Coding Conditions 67
 Character Names 56
 Cleanup, stack frame kind 28
 Code Coverage 127
 Compatibility with other Lisps 20
 Compile time type errors 23
 Compiler Diagnostic Severity 16
 Compiler messages 15
 Concurrency 121
 Converting between Strings and Octet Vectors 68

D

Deadline 64
 Debug optimization quality 30, 32
 Debug variables 29
 Debugger 26
 debugger, disabling 36
 debugger, enabling 36
 Decimal Syntax for Rationals 42
 declaration, **dynamic-extent** 37
 Declaration, **freeze-type** 66
 Declarations 92
 Deprecation Conditions 141
 Deprecation Declaration 142
 Deprecation Examples 142
 disabling debugger 36
 disabling ldb 12, 36
dynamic-extent declaration 37

E

Efficiency 37
 Entry points, external 28
 Errors, run-time 29
 Existing programs, to run 20
 Extended Package Prefix Syntax 42
 External entry points 28
 External Format Designators 67
 External formats 73
 External, stack frame kind 28

F

Fast Read Lock 125
 Finalization 44
 Foreign Function Interface, generation 129
freeze-type declaration 66
 Frlock 125
 Function, tracing 34

G

Garbage collection 44
 Garbage Collection, conservative 8
 Garbage Collection, generational 8
 Gate 124

H

Hash tables 60
 Hashing, cryptographic 132

I

Inline expansion 24, 32
 Interning Symbols 42
 Interpreter 25
 Interrupts 29
 Introspecting Deprecation Information 142

L

ldb 2, 12
 ldb, disabling 12, 36
 ldb, enabling 36
 Locations, unknown 29
 Logical pathnames 83

M

Macroexpansion 18
 Macroexpansion, errors during 24
 Mailbox, lock-free 123
 Messages, Compiler 15
 Modular arithmetic 39, 139

N

Name, of character 56
 NFKC 42
 Nicknames, Package-local 42
 Normalization, String 58
 Normalization, Symbol Name 42

O

Open-coding 24
 Operating System Interface 133
 optimization quality, **safety** 38
 Optimize declaration 32
 Optional, stack frame kind 28
 Original Source 17, 18

P

Package Locks	42
Package Prefix Syntax, extended	42
Package-Local Nicknames	42
Packages, locked	92
Pathnames	83
Pathnames, logical	83
Policy, debugger	32
Posix	133
Precise type checking	19
Processing Path	17, 18
Profiling	114
Profiling, deterministic	114
Profiling, statistical	114

Q

Queue, FIFO	138
Queue, lock-free	122

R

Random Number Generation	62
Rational, decimal syntax for	42
Read errors, compiler	24
Read-Eval-Print Loop	120
Reader Extensions	42
Recursion, tail	28
REPL	120

S

Safety optimization quality	19, 38
safety optimization quality	38
Sb-concurrency	121
Semi-inline expansion	32
Severity of compiler messages	16
Single Stepping	35
Slot access	37
Sockets, Networking	110
Source location printing, debugger	31

Source-to-source transformation	18
Stack frames	27
Static functions	24
Stepper	35
Stream External formats	85
Supported External Formats	68
Symbol Name Normalization	42
Symbols, interned	42
Synchronous Timeout	64
System Calls	133

T

Tail recursion	28
The Default External Format	67
The Deprecation Pipeline	140
Timeout	63, 64
Tracing	34
Type checking, at compile time	23
Type checking, precise	19
Types, portability	20

U

Unicode	42, 56
Unknown code locations	29

V

Validity of debug variables	30
Variables, debugger access	29

W

Weak pointers	45
Why Deprecate?	140

Appendix B Function Index

?

? 34

A

 abort 33
 sb-thread:abort-thread 97
 sb-ext:add-implementation-package 95
 sb-ext:add-package-local-nickname 43
 sb-alien:addr 74
sb-sequence:adjust-sequence 50
sb-unicode:age 56
 sb-alien:alien-callable-function 79
 sb-alien:alien-funcall 78
 sb-alien:alien-sap 74
sb-unicode:alphabetic-p 57
 sb-ext:always-bound 40
 cl:apropos 140
 sb-ext:array-storage-vector 65
 sb-ext:assert-version->= 66
 sb-ext:atomic-decf 99
 sb-ext:atomic-incf 99
 sb-ext:atomic-pop 100
 sb-ext:atomic-push 100
 sb-ext:atomic-update 100

B

 backtrace 34
 sb-thread:barrier 105
sb-unicode:bidirectional-class 56
sb-unicode:bidirectional-mirroring-glyph 56
 cl:both-case-p 58
 bottom 27
 sb-ext:bytes-consed-between-gcs 45

C

 sb-ext:cancel-finalization 45
 sb-ext:cas 101
sb-unicode:case-ignorable-p 57
sb-unicode:cased-p 57
sb-unicode:casefold 58
 sb-alien:cast 74
 cl:change-class 49
sb-unicode:char-block 57
 cl:char-downcase 58
 cl:char-name 56
 cl:class-name 49
 cl:class-of 48
 sb-mop:class-prototype 48
 sb-cover:clear-coverage 127
 sb-thread:clear-semaphore-notification 103
 cl:close 86
sb-concurrency:close-gate 124
 cl:coerce 49
sb-unicode:combining-class 56
 sb-ext:compare-and-swap 100
 sb-mop:compute-effective-method 47
sb-sequence:concatenate 51
 sb-thread:condition-broadcast 105
 sb-thread:condition-notify 105

 sb-thread:condition-wait 104
sb-unicode:confusable-p 59
 cl:cons 38
 continue 33
copy function 52
 cl:copy-seq 50

D

sb-unicode:decimal-value 56
 cl:declare 92
sb-unicode:default-ignorable-p 57
 cl:defclass 49
 cl:defconstant 5
 sb-ext:defglobal 39
 sb-alien:define-alien-callable 79
 sb-alien:define-alien-routine 78
 sb-alien:define-alien-variable 76
 sb-ext:define-hash-table-test 61
 cl:defmethod 49
 cl:defpackage 42, 92, 95
 cl:defstruct 49
 sb-ext:delete-directory 65
 sb-ext:deprecated 142
sb-concurrency:dequeue 122
 sb-alien:deref 74
 describe 34
 sb-ext:describe-compiler-policy 22
sb-unicode:digit-value 56
 sb-ext:disable-debugger 36
 sb-ext:disable-package-locks 92, 94
 cl:dolist 51
sb-sequence:dosequence 51
 down 27
 sb-ext:dynamic-space-size 45

E

 sb-unicode:east-asian-width 57
 cl:ed 59
element function 52
sb-sequence:elt 50
sb-sequence:empty 51
 sb-ext:enable-debugger 36
 sb-ext:enable-package-locks 92, 94
endp function 52
sb-concurrency:enqueue 122
 sb-mop:ensure-class 49
 sb-mop:ensure-class-using-class 49
 cl:ensure-generic-function 48
 error 34
 cl:eval 25
 sb-ext:exit 9
 cl:expt 62
 sb-alien:extern-alien 76

F

sb-posix:file-descriptor	133
sb-posix:filename	134
sb-ext:finalize	44
sb-mop:finalize-inheritance	48
cl:find	49
cl:find-class	49
cl:find-method	49
cl:find-package	42
cl:find-symbol	42
cl:flet	38, 92
sb-ext:float-denormalized-p	66
frame	27
sb-alien:free-alien	75
sb-concurrency:frlock-name	125
sb-concurrency:frlock-read	125
sb-concurrency:frlock-read-begin	125
sb-concurrency:frlock-read-end	125
sb-concurrency:frlock-write	125
sb-mop:funcallable-standard-instance-access	49

G

sb-concurrency:gate-name	124
sb-concurrency:gate-open-p	124
sb-concurrency:gatep	124
sb-ext:gc	44
sb-ext:gc-logfile	45
sb-unicode:general-category	56
sb-ext:generation-average-age	45
sb-ext:generation-bytes-allocated	46
sb-ext:generation-bytes-consed-between-gcs	46
sb-ext:generation-minimum-age-before-gc	46
sb-ext:generation-number-of-gcs	46
sb-ext:generation-number-of-gcs-before-promotion	46
sb-mop:generic-function-declarations	48
sb-ext:get-bytes-consed	45
sb-alien:get-errno	76
sb-bsd-sockets:get-host-by-address	113
sb-bsd-sockets:get-host-by-name	113
sb-bsd-sockets:get-protocol-by-name	113
sb-ext:get-time-of-day	66
sb-posix:getcwd	137
sb-ext:global	39
sb-concurrency:grab-frlock-write-lock	126
sb-thread:grab-mutex	102
sb-unicode:grapheme-break-class	57
sb-unicode:graphemes	59

H

sb-unicode:hangul-syllable-type	57
sb-ext:hash-table-synchronized-p	62
sb-ext:hash-table-weakness	62
help	34
sb-unicode:hex-digit-p	57
sb-bsd-sockets:host-ent-address	113

I

sb-unicode:ideographic-p	57
index function	52
cl:inspect	60
sb-sys:int-sap	74
cl:intern	42
sb-mop:intern-eql-specializer	49
sb-thread:interrupt-thread	97
sb-sequence:iterator-copy	53
sb-sequence:iterator-element	53
sb-sequence:iterator-endp	53
sb-sequence:iterator-index	53
sb-sequence:iterator-step	52

J

sb-thread:join-thread	63, 97
-----------------------	--------

L

cl:labels	38, 92
sb-sequence:length	50
cl:let	40, 92
cl:let*	40, 92
sb-unicode:line-break-class	58
sb-unicode:lines	59
cl:list	38
cl:list*	38
sb-thread:list-all-threads	96
sb-ext:list-all-timers	109
list-locals	29
sb-concurrency:list-mailbox-messages	123
sb-concurrency:list-queue-contents	122
sb-alien:load-shared-object	77
sb-ext:lock-package	95
cl:logand	39
cl:logical-pathname-translations	83
sb-unicode:lowercase	58
sb-unicode:lowercase-p	57

M

cl:macrolet	92
sb-concurrency:mailbox-count	123
sb-concurrency:mailbox-empty-p	123
sb-concurrency:mailbox-name	123
sb-concurrency:mailboxp	123
sb-thread:main-thread	96
sb-thread:main-thread-p	96
sb-alien:make-alien	75
sb-alien:make-alien-string	75
cl:make-array	38
sb-concurrency:make-frlock	125
sb-concurrency:make-gate	124
cl:make-hash-table	60
sb-bsd-sockets:make-inet-address	113
sb-bsd-sockets:make-inet6-address	113
cl:make-instance	49
sb-concurrency:make-mailbox	123
sb-mop:make-method-lambda	49
sb-pcl:make-method-specializers-form	49
sb-thread:make-mutex	102
sb-concurrency:make-queue	122
cl:make-random-state	62
sb-thread:make-semaphore	103

sb-thread:make-semaphore-notification 103
 sb-sequence:make-sequence-iterator 52
 sb-sequence:make-sequence-like 50
 sb-sequence:make-simple-sequence-iterator .. 52, 53
 sb-thread:make-thread 96
 sb-ext:make-timer 108
 sb-thread:make-waitqueue 104
 sb-ext:make-weak-pointer 45
 sb-sequence:map 51
 sb-sprof:map-traces 117
 sb-unicode:math-p 57
 sb-md5:md5sum-file 132
 sb-md5:md5sum-sequence 132
 sb-md5:md5sum-stream 132
 sb-md5:md5sum-string 132
 sb-sequence:merge 51
 sb-unicode:mirrored-p 56
 sb-ext:muffle-conditions 15
 sb-thread:mux-name 102
 sb-thread:mux-owner 102
 sb-thread:mux-value 102

N

cl:name-char 56
 sb-ext:name-conflict-symbols 60
 sb-ext:native-namestring 83
 sb-ext:native-pathname 83
 next 35
 sb-bsd-sockets:non-blocking-mode 111
 sb-unicode:normalize-string 58
 sb-unicode:normalized-p 58
 sb-unicode:numeric-value 56

O

sb-ext:octets-to-string 68
 cl:open 67, 85
 sb-concurrency:open-gate 124
 out 35

P

sb-ext:package-implemented-by-list 95
 sb-ext:package-implements-list 95
 sb-ext:package-local-nicknames 43
 sb-ext:package-locally-nicknamed-by-list 43
 sb-ext:package-locked-error-symbol 95
 sb-ext:package-locked-p 95
 sb-ext:parse-native-namestring 83
 sb-pcl:parse-specializer-using-class 49
 sb-ext:posix-getenv 53
 print 34
 sb-ext:process-alive-p 55
 sb-ext:process-close 55
 sb-ext:process-core-dumped 55
 sb-ext:process-error 55
 sb-ext:process-exit-code 55
 sb-ext:process-input 55
 sb-ext:process-kill 55
 sb-ext:process-output 55
 sb-ext:process-p 55
 sb-ext:process-status 55
 sb-ext:process-wait 55
 sb-profile:profile 114

sb-sprof:profile-call-counts 118
 cl:prog2 3
 sb-unicode:proplist-p 57
 sb-ext:purify 66

Q

sb-concurrency:queue-count 122
 sb-concurrency:queue-empty-p 122
 sb-concurrency:queue-name 122
 sb-concurrency:queuep 122

R

cl:random 62
 cl:read 42
 cl:read-from-string 42
 cl:read-line 64
 sb-posix:readlink 137
 sb-ext:readtable-normalization 42
 sb-concurrency:receive-message 123
 sb-concurrency:receive-message-no-hang 123
 sb-concurrency:receive-pending-messages 123
 sb-concurrency:release-frlock-write-lock 126
 sb-thread:release-mutex 102
 sb-ext:remove-implementation-package 95
 sb-ext:remove-package-local-nickname 43
 sb-profile:report 114, 117, 127
 cl:require 59
 sb-profile:reset 114, 117
 sb-cover:reset-coverage 127
 restart 33
 restart-frame 33
 sb-cover:restore-coverage 127
 sb-cover:restore-coverage-from-file 128
 sb-ext:restrict-compiler-policy 22
 return 33
 sb-thread:return-from-thread 96
 sb-rotate-byte:rotate-byte 132, 139
 sb-ext:run-program 53, 64

S

sb-sprof:sample-pc 117
 sb-alien:sap-alien 74
 sb-sys:sap-ref-32 74
 sb-sys:sap= 74
 cl:satisfies 19
 sb-cover:save-coverage 127
 sb-cover:save-coverage-in-file 127
 sb-ext:save-lisp-and-die 10
 sb-ext:schedule-timer 108
 sb-unicode:script 57
 sb-ext:search-roots 46
 sb-ext:seed-random-state 62
 sb-thread:semaphore-count 103
 sb-thread:semaphore-name 103
 sb-thread:semaphore-notification-status 103
 sb-concurrency:send-message 123
 sb-unicode:sentence-break-class 58
 sb-unicode:sentences 59
 sb-ext:set-sbcl-source-location 83
 cl:setf 40
 setf element function 52
 cl:setq 40
 sb-thread:signal-semaphore 103

cl:sleep	63, 64
sb-alien:slot	74
sb-mop:slot-boundp-using-class	48, 49
sb-mop:slot-definition-allocation	49
sb-mop:slot-definition-name	49
sb-mop:slot-value-using-class	48, 49
sb-bsd-sockets:socket-accept	110
sb-bsd-sockets:socket-bind	110
sb-bsd-sockets:socket-close	111
sb-bsd-sockets:socket-connect	110
sb-bsd-sockets:socket-error	111
sb-bsd-sockets:socket-listen	111
sb-bsd-sockets:socket-make-stream	111
sb-bsd-sockets:socket-name	110
sb-bsd-sockets:socket-open-p	111
sb-bsd-sockets:socket-peername	110
sb-bsd-sockets:socket-receive	110
sb-bsd-sockets:socket-send	111
sb-bsd-sockets:socket-shutdown	111
sb-bsd-sockets:sockopt-broadcast	112
sb-bsd-sockets:sockopt-bsd-compatible	112
sb-bsd-sockets:sockopt-debug	112
sb-bsd-sockets:sockopt-dont-route	112
sb-bsd-sockets:sockopt-keep-alive	112
sb-bsd-sockets:sockopt-oob-inline	112
sb-bsd-sockets:sockopt-pass-credentials	112
sb-bsd-sockets:sockopt-reuse-address	112
sb-bsd-sockets:sockopt-tcp-nodelay	112
sb-unicode:soft-dotted-p	57
source	31
sb-mop:standard-instance-access	49
start	35
sb-sprof:start-profiling	117
step	35, 36
step function	52
stop	35
sb-sprof:stop-profiling	118
sb-gray:stream-advance-to-column	87
sb-gray:stream-clear-input	86
sb-gray:stream-clear-output	87
cl:stream-element-type	86
cl:stream-external-format	85
sb-gray:stream-file-position	86
sb-gray:stream-finish-output	87
sb-gray:stream-force-output	87
sb-gray:stream-fresh-line	87
sb-gray:stream-line-column	88
sb-gray:stream-line-length	88
sb-gray:stream-listen	87
sb-gray:stream-peek-char	87
sb-gray:stream-read-byte	88
sb-gray:stream-read-char	87
sb-gray:stream-read-char-no-hang	87
sb-gray:stream-read-line	87
sb-gray:stream-read-sequence	86
sb-gray:stream-start-line-p	88
sb-gray:stream-terpri	88
sb-gray:stream-unread-char	87
sb-gray:stream-write-byte	88
sb-gray:stream-write-char	88
sb-gray:stream-write-sequence	87
sb-gray:stream-write-string	88
sb-ext:string-to-octets	68
cl:string-upcase	58
cl:subseq	49, 50
cl:subtypep	48

cl:symbol-macrolet	92
sb-thread:symbol-value-in-thread	98
sb-posix:syslog	137

T

sb-thread:terminate-thread	98
sb-thread:thread-alive-p	96
sb-thread:thread-error-thread	99
sb-thread:thread-name	96
sb-thread:thread-yield	97
sb-ext:timer-name	108
sb-ext:timer-scheduled-p	108
sb-unicode:titlecase	58
top	27
toplevel	33
cl:trace	34, 60
sb-ext:truly-the	66
sb-thread:try-semaphore	103
cl:typep	48

U

sb-unicode:unicode-1-name	57
sb-unicode:unicode-equal	58
sb-unicode:unicode<	58
sb-unicode:unicode<=	59
sb-unicode:unicode=	58
sb-unicode:unicode>	59
sb-unicode:unicode>=	59
sb-alien:unload-shared-object	77
sb-ext:unlock-package	95
sb-ext:unmuffle-conditions	15
sb-pcl:unparse-specializer-using-class	49
sb-profile:unprofile	114
sb-sprof:unprofile-call-counts	118
sb-ext:unschedule-timer	108
cl:untrace	35
up	27
sb-unicode:uppercase	58
sb-unicode:uppercase-p	57

V

sb-mop:validate-superclass	48, 49
sb-debug:var	29
cl:vector	38

W

sb-ext:wait-for	64
sb-concurrency:wait-on-gate	124
sb-thread:wait-on-semaphore	103
sb-thread:waitqueue-name	104
sb-ext:weak-pointer-value	45
sb-unicode:whitespace-p	57
sb-alien:with-alien	75
cl:with-compilation-unit	16, 22
sb-sys:with-deadline	64
sb-ext:with-locked-hash-table	62
sb-thread:with-mutex	101
cl:with-open-file	67, 85
sb-sprof:with-profiling	116
sb-thread:with-recursive-lock	102
sb-sprof:with-sampling	116
sb-sequence:with-sequence-iterator	52

sb-sequence:with-sequence-iterator-functions ..	52
sb-ext:with-timeout	64
sb-ext:with-unlocked-packages	95
sb-ext:without-package-locks	95
sb-unicode:word-break-class	57
sb-unicode:words	59

Appendix C Variable Index

+

sb-pcl:+slot-unbound+ 49

A

sb-ext:*after-gc-hooks* 44

C

sb-aclrepl:*command-char* 120

sb-ext:*compiler-print-variable-alist* 15

sb-ext:*core-pathname* 12

sb-thread:*current-thread* 96

D

sb-ext:*debug-print-variable-alist* 27

E

sb-ext:*ed-functions* 59

sb-ext:*evaluator-mode* 25

sb-ext:*exit-hooks* 14

sb-aclrepl:*exit-on-eof* 120

F

cl:*features* 42

G

sb-ext:*gc-run-time* 45

I

sb-ext:*init-hooks* 14

sb-ext:*invoke-debugger-hook* 26

M

sb-aclrepl:*max-history* 120

sb-sprof:*max-samples* 118

sb-debug:*max-trace-indentation* 35

sb-ext:*module-provider-functions* 59

sb-ext:*muffled-warnings* 60

O

sb-ext:*on-package-variance* 44

P

cl:*package* 42, 92

sb-ext:*posix-argv* 9, 53

sb-aclrepl:*prompt* 120

R

cl:*random-state* 62

cl:*read-default-float-format* 42

S

sb-sprof:*sample-interval* 118

sb-ext:*save-hooks* 11

sb-ext:*stack-allocate-dynamic-extent* 37

sb-ext:*sysinit-pathname-function* 11

T

sb-debug:*trace-encapsulate-default* 35

sb-debug:*trace-indentation-step* 35

U

sb-aclrepl:*use-short-package-name* 120

sb-ext:*userinit-pathname-function* 11

Appendix D Type Index

B

cl:built-in-class 49

C

cl:character 3
 sb-ext:code-deletion-note 16
 sb-ext:compiler-note 16

D

sb-sys:deadline-timeout 64
 sb-ext:deprecation-condition 141
 sb-ext:deprecation-error 141, 142

E

sb-ext:early-deprecation-warning 141, 143
 cl:error 16, 141

F

sb-posix:file-descriptor 133
 sb-posix:file-descriptor-designator 133
 sb-posix:filename 133
 sb-posix:filename-designator 134
 sb-ext:final-deprecation-warning 141, 145
 cl:float 42, 62
 sb-posix:flock 134
 sb-concurrency:frlock 125
 sb-mop:funcallable-standard-class 48
 sb-mop:funcallable-standard-object 48
 cl:function 48
 sb-gray:fundamental-binary-input-stream 86
 sb-gray:fundamental-binary-output-stream 86
 sb-gray:fundamental-binary-stream 86
 sb-gray:fundamental-character-input-stream 86
 sb-gray:fundamental-character-output-stream 86
 sb-gray:fundamental-character-stream 86
 sb-gray:fundamental-input-stream 85
 sb-gray:fundamental-output-stream 86
 sb-gray:fundamental-stream 85

G

sb-concurrency:gate 124
 cl:generic-function 48

H

sb-bsd-sockets:host-ent 113

I

sb-bsd-sockets:inet-socket 112
 sb-bsd-sockets:inet6-socket 112
 sb-thread:interrupt-thread-error 99

J

sb-thread:join-thread-error 63, 99

L

sb-ext:late-deprecation-warning 141, 144
 cl:list 49
 sb-bsd-sockets:local-abstract-socket 113
 sb-bsd-sockets:local-socket 113

M

sb-concurrency:mailbox 123
 sb-thread:mutex 101

N

sb-ext:name-conflict 60
 cl:nil 3

P

cl:package-error 92
 sb-ext:package-lock-violation 92, 94
 sb-ext:package-locked-error 92, 94
 sb-posix:passwd 135
 sb-sequence:protocol-unimplemented 50

Q

sb-concurrency:queue 122

R

cl:rational 42

S

sb-thread:semaphore 103
 sb-thread:semaphore-notification 103
 cl:sequence 49, 50, 52
 sb-bsd-sockets:socket 110
 cl:standard-class 48
 cl:standard-generic-function 48
 cl:standard-object 48, 50
 sb-posix:stat 135
 cl:string 3
 cl:structure-class 49
 cl:style-warning 16, 141
 sb-ext:symbol-package-locked-error 92, 94

T

cl:t 49
 sb-posix:termios 136
 sb-thread:thread 96
 sb-thread:thread-error 98
 sb-ext:timeout 65
 sb-ext:timer 108
 sb-posix:timeval 136

V

`cl:vector` 49

W

`sb-thread:waitqueue` 104
`cl:warning` 16, 141

Colophon

This manual is maintained in Texinfo, and automatically translated into other forms (e.g. HTML or pdf). If you're *reading* this manual in one of these non-Texinfo translated forms, that's fine, but if you want to *modify* this manual, you are strongly advised to seek out a Texinfo version and modify that instead of modifying a translated version. Even better might be to seek out *the* Texinfo version (maintained at the time of this writing as part of the SBCL project at <http://sbcl.sourceforge.net/>) and submit a patch.