

NAME

perlsub - Perl subroutines

SYNOPSIS

To declare subroutines:

```
sub NAME;          # A "forward" declaration.
sub NAME(PROTO);   # ditto, but with prototypes
sub NAME : ATTRS;  # with attributes
sub NAME(PROTO) : ATTRS;  # with attributes and prototypes

sub NAME BLOCK     # A declaration and a definition.
sub NAME(PROTO) BLOCK  # ditto, but with prototypes
sub NAME(SIG) BLOCK   # with a signature instead
sub NAME : ATTRS BLOCK  # with attributes
sub NAME(PROTO) : ATTRS BLOCK # with prototypes and attributes
sub NAME(SIG) : ATTRS BLOCK  # with a signature and attributes
```

To define an anonymous subroutine at runtime:

```
$subref = sub BLOCK;    # no proto
$subref = sub (PROTO) BLOCK; # with proto
$subref = sub (SIG) BLOCK;    # with signature
$subref = sub : ATTRS BLOCK;  # with attributes
$subref = sub (PROTO) : ATTRS BLOCK; # with proto and attributes
$subref = sub (SIG) : ATTRS BLOCK;  # with signature and attributes
```

To import subroutines:

```
use MODULE qw(NAME1 NAME2 NAME3);
```

To call subroutines:

```
NAME(LIST);    # & is optional with parentheses.
NAME LIST;     # Parentheses optional if predeclared/imported.
&NAME(LIST);   # Circumvent prototypes.
&NAME;        # Makes current @_ visible to called subroutine.
```

DESCRIPTION

Like many languages, Perl provides for user-defined subroutines. These may be located anywhere in the main program, loaded in from other files via the `do`, `require`, or `use` keywords, or generated on the fly using `eval` or anonymous subroutines. You can even call a function indirectly using a variable containing its name or a CODE reference.

The Perl model for function call and return values is simple: all functions are passed as parameters one single flat list of scalars, and all functions likewise return to their caller one single flat list of scalars. Any arrays or hashes in these call and return lists will collapse, losing their identities--but you may always use pass-by-reference instead to avoid this. Both call and return lists may contain as many or as few scalar elements as you'd like. (Often a function without an explicit return statement is called a subroutine, but there's really no difference from Perl's perspective.)

Any arguments passed in show up in the array `@_`. (They may also show up in lexical variables introduced by a signature; see *Signatures* below.) Therefore, if you called a function with two arguments, those would be stored in `$_[0]` and `$_[1]`. The array `@_` is a local array, but its elements are aliases for the actual scalar parameters. In particular, if an element `$_[0]` is updated, the corresponding argument is updated (or an error occurs if it is not updatable). If an argument is an

array or hash element which did not exist when the function was called, that element is created only when (and if) it is modified or a reference to it is taken. (Some earlier versions of Perl created the element whether or not the element was assigned to.) Assigning to the whole array `@_` removes that aliasing, and does not update any arguments.

A `return` statement may be used to exit a subroutine, optionally specifying the returned value, which will be evaluated in the appropriate context (list, scalar, or void) depending on the context of the subroutine call. If you specify no return value, the subroutine returns an empty list in list context, the undefined value in scalar context, or nothing in void context. If you return one or more aggregates (arrays and hashes), these will be flattened together into one large indistinguishable list.

If no `return` is found and if the last statement is an expression, its value is returned. If the last statement is a loop control structure like a `foreach` or a `while`, the returned value is unspecified. The empty sub returns the empty list.

Aside from an experimental facility (see *Signatures* below), Perl does not have named formal parameters. In practice all you do is assign to a `my ()` list of these. Variables that aren't declared to be private are global variables. For gory details on creating private variables, see *Private Variables via `my()`* and *Temporary Values via `local()`*. To create protected environments for a set of functions in a separate package (and probably a separate file), see *"Packages" in `perlmod`*.

Example:

```
sub max {
my $max = shift(@_);
foreach $foo (@_) {
    $max = $foo if $max < $foo;
}
return $max;
}
$bestday = max($mon,$tue,$wed,$thu,$fri);
```

Example:

```
# get a line, combining continuation lines
# that start with whitespace

sub get_line {
    $thisline = $lookahead; # global variables!
    LINE: while (defined($lookahead = <STDIN>)) {
        if ($lookahead =~ /^[ \t]/) {
            $thisline .= $lookahead;
        }
        else {
            last LINE;
        }
    }
    return $thisline;
}

$lookahead = <STDIN>; # get first line
while (defined($line = get_line())) {
    ...
}
```

Assigning to a list of private variables to name your arguments:

```
sub maybeaset {
    my($key, $value) = @_;
    $Foo{$key} = $value unless $Foo{$key};
}
```

Because the assignment copies the values, this also has the effect of turning call-by-reference into call-by-value. Otherwise a function is free to do in-place modifications of `@_` and change its caller's values.

```
upcase_in($v1, $v2); # this changes $v1 and $v2
sub upcase_in {
    for (@_) { tr/a-z/A-Z/ }
}
```

You aren't allowed to modify constants in this way, of course. If an argument were actually literal and you tried to change it, you'd take a (presumably fatal) exception. For example, this won't work:

```
upcase_in("frederick");
```

It would be much safer if the `upcase_in()` function were written to return a copy of its parameters instead of changing them in place:

```
($v3, $v4) = upcase($v1, $v2); # this doesn't change $v1 and $v2
sub upcase {
    return unless defined wantarray; # void context, do nothing
    my @parms = @_;
    for (@parms) { tr/a-z/A-Z/ }
    return wantarray ? @parms : $parms[0];
}
```

Notice how this (unprototyped) function doesn't care whether it was passed real scalars or arrays. Perl sees all arguments as one big, long, flat parameter list in `@_`. This is one area where Perl's simple argument-passing style shines. The `upcase()` function would work perfectly well without changing the `upcase()` definition even if we fed it things like this:

```
@newlist = upcase(@list1, @list2);
@newlist = upcase( split /\:/, $var );
```

Do not, however, be tempted to do this:

```
(@a, @b) = upcase(@list1, @list2);
```

Like the flattened incoming parameter list, the return list is also flattened on return. So all you have managed to do here is stored everything in `@a` and made `@b` empty. See *Pass by Reference* for alternatives.

A subroutine may be called using an explicit `&` prefix. The `&` is optional in modern Perl, as are parentheses if the subroutine has been predeclared. The `&` is *not* optional when just naming the subroutine, such as when it's used as an argument to `defined()` or `undef()`. Nor is it optional when you want to do an indirect subroutine call with a subroutine name or reference using the `&$subref()` or `&{$subref}()` constructs, although the `$subref->()` notation solves that problem. See *perlref* for more about all that.

Subroutines may be called recursively. If a subroutine is called using the `&` form, the argument list is optional, and if omitted, no `@_` array is set up for the subroutine: the `@_` array at the time of the call is visible to subroutine instead. This is an efficiency mechanism that new users may wish to avoid.

```
&foo(1,2,3); # pass three arguments
foo(1,2,3);  # the same

foo();      # pass a null list
&foo();     # the same

&foo;      # foo() get current args, like foo(@_) !!
foo;       # like foo() IFF sub foo predeclared, else "foo"
```

Not only does the `&` form make the argument list optional, it also disables any prototype checking on arguments you do provide. This is partly for historical reasons, and partly for having a convenient way to cheat if you know what you're doing. See *Prototypes* below.

Since Perl 5.16.0, the `__SUB__` token is available under `use` feature `'current_sub'` and `use 5.16.0`. It will evaluate to a reference to the currently-running sub, which allows for recursive calls without knowing your subroutine's name.

```
use 5.16.0;
my $factorial = sub {
    my ($x) = @_;
    return 1 if $x == 1;
    return($x * __SUB__->( $x - 1 ) );
};
```

The behavior of `__SUB__` within a regex code block (such as `/ (? { . . }) /`) is subject to change.

Subroutines whose names are in all upper case are reserved to the Perl core, as are modules whose names are in all lower case. A subroutine in all capitals is a loosely-held convention meaning it will be called indirectly by the run-time system itself, usually due to a triggered event. Subroutines whose name start with a left parenthesis are also reserved the same way. The following is a list of some subroutines that currently do special, pre-defined things.

documented later in this document

AUTOLOAD

documented in *perlmod*

CLONE, CLONE_SKIP,

documented in *perlobj*

DESTROY

documented in *perlty*

BINMODE, CLEAR, CLOSE, DELETE, DESTROY, EOF, EXISTS, EXTEND, FETCH, FETCHSIZE, FILENO, FIRSTKEY, GETC, NEXTKEY, OPEN, POP, PRINT, PRINTF, PUSH, READ, READLINE, SCALAR, SEEK, SHIFT, SPLICE, STORE, STORESIZE, TELL, TIEARRAY, TIEHANDLE, TIEHASH, TIESCALAR, UNSHIFT, UNTIE, WRITE

documented in *PerlIO::via*

BINMODE, CLEARERR, CLOSE, EOF, ERROR, FDOPEN, FILENO, FILL, FLUSH, OPEN, POPPED, PUSHED, READ, SEEK, SETLINEBUF, SYSOPEN, TELL, UNREAD, UTF8, WRITE

documented in *perlfunc*

import, unimport, INC

documented in *UNIVERSAL*

VERSION

documented in *perldebguts*

DB::DB, DB::sub, DB::lsub, DB::goto, DB::postponed

undocumented, used internally by the *overload* feature

any starting with (

The `BEGIN`, `UNITCHECK`, `CHECK`, `INIT` and `END` subroutines are not so much subroutines as named special code blocks, of which you can have more than one in a package, and which you can **not** call explicitly. See "*BEGIN, UNITCHECK, CHECK, INIT and END*" in *perlmod*

Signatures

WARNING: Subroutine signatures are experimental. The feature may be modified or removed in future versions of Perl.

Perl has an experimental facility to allow a subroutine's formal parameters to be introduced by special syntax, separate from the procedural code of the subroutine body. The formal parameter list is known as a *signature*. The facility must be enabled first by a pragmatic declaration, use `feature 'signatures'`, and it will produce a warning unless the "experimental::signatures" warnings category is disabled.

The signature is part of a subroutine's body. Normally the body of a subroutine is simply a braced block of code. When using a signature, the signature is a parenthesised list that goes immediately after the subroutine name. The signature declares lexical variables that are in scope for the block. When the subroutine is called, the signature takes control first. It populates the signature variables from the list of arguments that were passed. If the argument list doesn't meet the requirements of the signature, then it will throw an exception. When the signature processing is complete, control passes to the block.

Positional parameters are handled by simply naming scalar variables in the signature. For example,

```
sub foo ($left, $right) {  
    return $left + $right;  
}
```

takes two positional parameters, which must be filled at runtime by two arguments. By default the parameters are mandatory, and it is not permitted to pass more arguments than expected. So the above is equivalent to

```
sub foo {  
    die "Too many arguments for subroutine" unless @_ <= 2;  
    die "Too few arguments for subroutine" unless @_ >= 2;  
    my $left = $_[0];  
    my $right = $_[1];  
    return $left + $right;  
}
```

An argument can be ignored by omitting the main part of the name from a parameter declaration, leaving just a bare `$` sigil. For example,

```
sub foo ($first, $, $third) {  
    return "first=$first, third=$third";  
}
```

Although the ignored argument doesn't go into a variable, it is still mandatory for the caller to pass it.

A positional parameter is made optional by giving a default value, separated from the parameter name by `=:`

```
sub foo ($left, $right = 0) {  
    return $left + $right;  
}
```

The above subroutine may be called with either one or two arguments. The default value expression is evaluated when the subroutine is called, so it may provide different default values for different calls. It is only evaluated if the argument was actually omitted from the call. For example,

```
my $auto_id = 0;  
sub foo ($thing, $id = $auto_id++) {  
    print "$thing has ID $id";  
}
```

automatically assigns distinct sequential IDs to things for which no ID was supplied by the caller. A default value expression may also refer to parameters earlier in the signature, making the default for one parameter vary according to the earlier parameters. For example,

```
sub foo ($first_name, $surname, $nickname = $first_name) {  
    print "$first_name $surname is known as \"$nickname\"";  
}
```

An optional parameter can be nameless just like a mandatory parameter. For example,

```
sub foo ($thing, $ = 1) {  
    print $thing;  
}
```

The parameter's default value will still be evaluated if the corresponding argument isn't supplied, even though the value won't be stored anywhere. This is in case evaluating it has important side effects. However, it will be evaluated in void context, so if it doesn't have side effects and is not trivial it will generate a warning if the "void" warning category is enabled. If a nameless optional parameter's default value is not important, it may be omitted just as the parameter's name was:

```
sub foo ($thing, $=) {  
    print $thing;  
}
```

Optional positional parameters must come after all mandatory positional parameters. (If there are no mandatory positional parameters then an optional positional parameters can be the first thing in the signature.) If there are multiple optional positional parameters and not enough arguments are supplied to fill them all, they will be filled from left to right.

After positional parameters, additional arguments may be captured in a slurpy parameter. The simplest form of this is just an array variable:

```
sub foo ($filter, @inputs) {  
    print $filter->($_) foreach @inputs;  
}
```

With a slurpy parameter in the signature, there is no upper limit on how many arguments may be passed. A slurpy array parameter may be nameless just like a positional parameter, in which case its only effect is to turn off the argument limit that would otherwise apply:

```
sub foo ($thing, @) {  
    print $thing;  
}
```

A slurpy parameter may instead be a hash, in which case the arguments available to it are interpreted as alternating keys and values. There must be as many keys as values: if there is an odd argument then an exception will be thrown. Keys will be stringified, and if there are duplicates then the later instance takes precedence over the earlier, as with standard hash construction.

```
sub foo ($filter, %inputs) {  
    print $filter->($_, $inputs{$_}) foreach sort keys %inputs;  
}
```

A slurpy hash parameter may be nameless just like other kinds of parameter. It still insists that the number of arguments available to it be even, even though they're not being put into a variable.

```
sub foo ($thing, %) {  
    print $thing;  
}
```

A slurpy parameter, either array or hash, must be the last thing in the signature. It may follow mandatory and optional positional parameters; it may also be the only thing in the signature. Slurpy parameters cannot have default values: if no arguments are supplied for them then you get an empty array or empty hash.

A signature may be entirely empty, in which case all it does is check that the caller passed no arguments:

```
sub foo () {  
    return 123;  
}
```

When using a signature, the arguments are still available in the special array variable `@_`, in addition to the lexical variables of the signature. There is a difference between the two ways of accessing the arguments: `@_` *aliases* the arguments, but the signature variables get *copies* of the arguments. So writing to a signature variable only changes that variable, and has no effect on the caller's variables, but writing to an element of `@_` modifies whatever the caller used to supply that argument.

There is a potential syntactic ambiguity between signatures and prototypes (see *Prototypes*), because both start with an opening parenthesis and both can appear in some of the same places, such as just after the name in a subroutine declaration. For historical reasons, when signatures are not enabled, any opening parenthesis in such a context will trigger very forgiving prototype parsing. Most signatures will be interpreted as prototypes in those circumstances, but won't be valid prototypes. (A valid prototype cannot contain any alphabetic character.) This will lead to somewhat confusing error messages.

To avoid ambiguity, when signatures are enabled the special syntax for prototypes is disabled. There is no attempt to guess whether a parenthesised group was intended to be a prototype or a signature. To give a subroutine a prototype under these circumstances, use a *prototype attribute*. For example,

```
sub foo :prototype($) { $_[0] }
```

It is entirely possible for a subroutine to have both a prototype and a signature. They do different jobs: the prototype affects compilation of calls to the subroutine, and the signature puts argument values into lexical variables at runtime. You can therefore write

```
sub foo ($left, $right) : prototype($$) {  
    return $left + $right;  
}
```

The prototype attribute, and any other attributes, come after the signature.

Private Variables via my()

Synopsis:

```
my $foo;          # declare $foo lexically local
my (@wid, %get);  # declare list of variables local
my $foo = "flurp"; # declare $foo lexical, and init it
my @oof = @bar;   # declare @oof lexical, and init it
my $x : Foo = $y; # similar, with an attribute applied
```

WARNING: The use of attribute lists on `my` declarations is still evolving. The current semantics and interface are subject to change. See *attributes* and *Attribute::Handlers*.

The `my` operator declares the listed variables to be lexically confined to the enclosing block, conditional (`if/unless/elsif/else`), loop (`for/foreach/while/until/continue`), subroutine, `eval`, or `do/require/use'd` file. If more than one value is listed, the list must be placed in parentheses. All listed elements must be legal lvalues. Only alphanumeric identifiers may be lexically scoped--magical built-ins like `$/` must currently be localized with `local` instead.

Unlike dynamic variables created by the `local` operator, lexical variables declared with `my` are totally hidden from the outside world, including any called subroutines. This is true if it's the same subroutine called from itself or elsewhere--every call gets its own copy.

This doesn't mean that a `my` variable declared in a statically enclosing lexical scope would be invisible. Only dynamic scopes are cut off. For example, the `bumpx()` function below has access to the lexical `$x` variable because both the `my` and the `sub` occurred at the same scope, presumably file scope.

```
my $x = 10;
sub bumpx { $x++ }
```

An `eval()`, however, can see lexical variables of the scope it is being evaluated in, so long as the names aren't hidden by declarations within the `eval()` itself. See *perlref*.

The parameter list to `my()` may be assigned to if desired, which allows you to initialize your variables. (If no initializer is given for a particular variable, it is created with the undefined value.) Commonly this is used to name input parameters to a subroutine. Examples:

```
$arg = "fred";    # "global" variable
$n = cube_root(27);
print "$arg thinks the root is $n\n";
fred thinks the root is 3
```

```
sub cube_root {
my $arg = shift; # name doesn't matter
$arg **= 1/3;
return $arg;
}
```

The `my` is simply a modifier on something you might assign to. So when you do assign to variables in its argument list, `my` doesn't change whether those variables are viewed as a scalar or an array. So

```
my ($foo) = <STDIN>; # WRONG?
my @FOO = <STDIN>;
```

both supply a list context to the right-hand side, while

```
my $foo = <STDIN>;
```


supplies a scalar context. But the following declares only one variable:

```
my $foo, $bar = 1;    # WRONG
```

That has the same effect as

```
my $foo;
$bar = 1;
```

The declared variable is not introduced (is not visible) until after the current statement. Thus,

```
my $x = $x;
```

can be used to initialize a new `$x` with the value of the old `$x`, and the expression

```
my $x = 123 and $x == 123
```

is false unless the old `$x` happened to have the value 123.

Lexical scopes of control structures are not bounded precisely by the braces that delimit their controlled blocks; control expressions are part of that scope, too. Thus in the loop

```
while (my $line = <>) {
    $line = lc $line;
} continue {
    print $line;
}
```

the scope of `$line` extends from its declaration throughout the rest of the loop construct (including the `continue` clause), but not beyond it. Similarly, in the conditional

```
if ((my $answer = <STDIN>) =~ /^yes$/i) {
    user_agrees();
} elsif ($answer =~ /^no$/i) {
    user_disagrees();
} else {
    chomp $answer;
    die "'$answer' is neither 'yes' nor 'no'";
}
```

the scope of `$answer` extends from its declaration through the rest of that conditional, including any `elsif` and `else` clauses, but not beyond it. See *"Simple Statements" in perlsyn* for information on the scope of variables in statements with modifiers.

The `foreach` loop defaults to scoping its index variable dynamically in the manner of `local`. However, if the index variable is prefixed with the keyword `my`, or if there is already a lexical by that name in scope, then a new lexical is created instead. Thus in the loop

```
for my $i (1, 2, 3) {
    some_function();
}
```

the scope of `$i` extends to the end of the loop, but not beyond it, rendering the value of `$i` inaccessible within `some_function()`.

Some users may wish to encourage the use of lexically scoped variables. As an aid to catching implicit uses to package variables, which are always global, if you say

```
use strict 'vars';
```

then any variable mentioned from there to the end of the enclosing block must either refer to a lexical variable, be predeclared via `our` or `use vars`, or else must be fully qualified with the package name. A compilation error results otherwise. An inner block may countermand this with `no strict 'vars'`.

A `my` has both a compile-time and a run-time effect. At compile time, the compiler takes notice of it. The principal usefulness of this is to quiet `use strict 'vars'`, but it is also essential for generation of closures as detailed in *perlref*. Actual initialization is delayed until run time, though, so it gets executed at the appropriate time, such as each time through a loop, for example.

Variables declared with `my` are not part of any package and are therefore never fully qualified with the package name. In particular, you're not allowed to try to make a package variable (or other global) lexical:

```
my $pack::var; # ERROR! Illegal syntax
```

In fact, a dynamic variable (also known as package or global variables) are still accessible using the fully qualified `::` notation even while a lexical of the same name is also visible:

```
package main;
local $x = 10;
my $x = 20;
print "$x and $::x\n";
```

That will print out 20 and 10.

You may declare `my` variables at the outermost scope of a file to hide any such identifiers from the world outside that file. This is similar in spirit to C's static variables when they are used at the file level. To do this with a subroutine requires the use of a closure (an anonymous function that accesses enclosing lexicals). If you want to create a private subroutine that cannot be called from outside that block, it can declare a lexical variable containing an anonymous sub reference:

```
my $secret_version = '1.001-beta';
my $secret_sub = sub { print $secret_version };
&$secret_sub();
```

As long as the reference is never returned by any function within the module, no outside module can see the subroutine, because its name is not in any package's symbol table. Remember that it's not *REALLY* called `$some_pack::secret_version` or anything; it's just `$secret_version`, unqualified and unqualifiable.

This does not work with object methods, however; all object methods have to be in the symbol table of some package to be found. See *"Function Templates" in perlref* for something of a work-around to this.

Persistent Private Variables

There are two ways to build persistent private variables in Perl 5.10. First, you can simply use the `state` feature. Or, you can use closures, if you want to stay compatible with releases older than 5.10.

Persistent variables via `state()`

Beginning with Perl 5.10.0, you can declare variables with the `state` keyword in place of `my`. For that to work, though, you must have enabled that feature beforehand, either by using the `feature` pragma, or by using `-E` on one-liners (see *feature*). Beginning with Perl 5.16, the `CORE::state` form does not require the `feature` pragma.

The `state` keyword creates a lexical variable (following the same scoping rules as `my`) that persists from one subroutine call to the next. If a state variable resides inside an anonymous subroutine, then each copy of the subroutine has its own copy of the state variable. However, the value of the state variable will still persist between calls to the same copy of the anonymous subroutine. (Don't forget that `sub { ... }` creates a new subroutine each time it is executed.)

For example, the following code maintains a private counter, incremented each time the `gimme_another()` function is called:

```
use feature 'state';
sub gimme_another { state $x; return ++$x }
```

And this example uses anonymous subroutines to create separate counters:

```
use feature 'state';
sub create_counter {
return sub { state $x; return ++$x }
}
```

Also, since `$x` is lexical, it can't be reached or modified by any Perl code outside.

When combined with variable declaration, simple scalar assignment to `state` variables (as in `state $x = 42`) is executed only the first time. When such statements are evaluated subsequent times, the assignment is ignored. The behavior of this sort of assignment to non-scalar variables is undefined.

Persistent variables with closures

Just because a lexical variable is lexically (also called statically) scoped to its enclosing block, `eval`, or `do FILE`, this doesn't mean that within a function it works like a C static. It normally works more like a C auto, but with implicit garbage collection.

Unlike local variables in C or C++, Perl's lexical variables don't necessarily get recycled just because their scope has exited. If something more permanent is still aware of the lexical, it will stick around. So long as something else references a lexical, that lexical won't be freed--which is as it should be. You wouldn't want memory being free until you were done using it, or kept around once you were done. Automatic garbage collection takes care of this for you.

This means that you can pass back or save away references to lexical variables, whereas to return a pointer to a C auto is a grave error. It also gives us a way to simulate C's function statics. Here's a mechanism for giving a function private variables with both lexical scoping and a static lifetime. If you do want to create something like C's static variables, just enclose the whole function in an extra block, and put the static variable outside the function but in the block.

```
{
my $secret_val = 0;
sub gimme_another {
    return ++$secret_val;
}
}
# $secret_val now becomes unreachable by the outside
# world, but retains its value between calls to gimme_another
```

If this function is being sourced in from a separate file via `require` or `use`, then this is probably just fine. If it's all in the main program, you'll need to arrange for the `my` to be executed early, either by putting the whole block above your main program, or more likely, placing merely a `BEGIN` code block around it to make sure it gets executed before your program starts to run:

```
BEGIN {
```

```
my $secret_val = 0;
sub gimme_another {
    return ++$secret_val;
}
}
```

See "*BEGIN, UNITCHECK, CHECK, INIT and END*" in *perlmod* about the special triggered code blocks, BEGIN, UNITCHECK, CHECK, INIT and END.

If declared at the outermost scope (the file scope), then lexicals work somewhat like C's file statics. They are available to all functions in that same file declared below them, but are inaccessible from outside that file. This strategy is sometimes used in modules to create private variables that the whole module can see.

Temporary Values via local()

WARNING: In general, you should be using `my` instead of `local`, because it's faster and safer. Exceptions to this include the global punctuation variables, global filehandles and formats, and direct manipulation of the Perl symbol table itself. `local` is mostly used when the current value of a variable must be visible to called subroutines.

Synopsis:

```
# localization of values

local $foo;          # make $foo dynamically local
local (@wid, %get);  # make list of variables local
local $foo = "flurp"; # make $foo dynamic, and init it
local @oof = @bar;   # make @oof dynamic, and init it

local $hash{key} = "val"; # sets a local value for this hash entry
delete local $hash{key}; # delete this entry for the current block
local ($cond ? $v1 : $v2); # several types of lvalues support
    # localization

# localization of symbols

local *FH;          # localize $FH, @FH, %FH, &FH ...
local *merlyn = *randal; # now $merlyn is really $randal, plus
                        # @merlyn is really @randal, etc
local *merlyn = 'randal'; # SAME THING: promote 'randal' to *randal
local *merlyn = \$randal; # just alias $merlyn, not @merlyn etc
```

A `local` modifies its listed variables to be "local" to the enclosing block, `eval`, or `do FILE--` and to *any subroutine called from within that block*. A `local` just gives temporary values to global (meaning package) variables. It does *not* create a local variable. This is known as dynamic scoping. Lexical scoping is done with `my`, which works more like C's auto declarations.

Some types of lvalues can be localized as well: hash and array elements and slices, conditionals (provided that their result is always localizable), and symbolic references. As for simple variables, this creates new, dynamically scoped values.

If more than one variable or expression is given to `local`, they must be placed in parentheses. This operator works by saving the current values of those variables in its argument list on a hidden stack and restoring them upon exiting the block, subroutine, or `eval`. This means that called subroutines can also reference the local variable, but not the global one. The argument list may be assigned to if desired, which allows you to initialize your local variables. (If no initializer is given for a particular

variable, it is created with an undefined value.)

Because `local` is a run-time operator, it gets executed each time through a loop. Consequently, it's more efficient to localize your variables outside the loop.

Grammatical note on `local()`

A `local` is simply a modifier on an lvalue expression. When you assign to a localized variable, the `local` doesn't change whether its list is viewed as a scalar or an array. So

```
local($foo) = <STDIN>;
local @FOO = <STDIN>;
```

both supply a list context to the right-hand side, while

```
local $foo = <STDIN>;
```

supplies a scalar context.

Localization of special variables

If you localize a special variable, you'll be giving a new value to it, but its magic won't go away. That means that all side-effects related to this magic still work with the localized value.

This feature allows code like this to work :

```
# Read the whole contents of FILE in $slurp
{ local $/ = undef; $slurp = <FILE>; }
```

Note, however, that this restricts localization of some values ; for example, the following statement dies, as of perl 5.10.0, with an error *Modification of a read-only value attempted*, because the `$!` variable is magical and read-only :

```
local $! = 2;
```

One exception is the default scalar variable: starting with perl 5.14 `local($_)` will always strip all magic from `$_`, to make it possible to safely reuse `$_` in a subroutine.

WARNING: Localization of tied arrays and hashes does not currently work as described. This will be fixed in a future release of Perl; in the meantime, avoid code that relies on any particular behavior of localising tied arrays or hashes (localising individual elements is still okay). See "*Localising Tied Arrays and Hashes Is Broken*" in *perl58delta* for more details.

Localization of globs

The construct

```
local *name;
```

creates a whole new symbol table entry for the glob `name` in the current package. That means that all variables in its glob slot (`$name`, `@name`, `%name`, `&name`, and the `name` filehandle) are dynamically reset.

This implies, among other things, that any magic eventually carried by those variables is locally lost. In other words, saying `local */` will not have any effect on the internal value of the input record separator.

Localization of elements of composite types

It's also worth taking a moment to explain what happens when you `localize` a member of a composite type (i.e. an array or hash element). In this case, the element is *localized by name*. This

means that when the scope of the `local()` ends, the saved value will be restored to the hash element whose key was named in the `local()`, or the array element whose index was named in the `local()`. If that element was deleted while the `local()` was in effect (e.g. by a `delete()` from a hash or a `shift()` of an array), it will spring back into existence, possibly extending an array and filling in the skipped elements with `undef`. For instance, if you say

```
%hash = ( 'This' => 'is', 'a' => 'test' );
@ary  = ( 0..5 );
{
    local($ary[5]) = 6;
    local($hash{'a'}) = 'drill';
    while (my $e = pop(@ary)) {
        print "$e . . .\n";
        last unless $e > 3;
    }
    if (@ary) {
        $hash{'only a'} = 'test';
        delete $hash{'a'};
    }
}
print join(' ', map { "$_ $hash{$_}" } sort keys %hash), ".\n";
print "The array has ", scalar(@ary), " elements: ",
      join(' ', map { defined $_ ? $_ : 'undef' } @ary), "\n";
```

Perl will print

```
6 . . .
4 . . .
3 . . .
This is a test only a test.
The array has 6 elements: 0, 1, 2, undef, undef, 5
```

The behavior of `local()` on non-existent members of composite types is subject to change in future.

Localized deletion of elements of composite types

You can use the `delete local $array[$idx]` and `delete local $hash{key}` constructs to delete a composite type entry for the current block and restore it when it ends. They return the array/hash value before the localization, which means that they are respectively equivalent to

```
do {
    my $val = $array[$idx];
    local $array[$idx];
    delete $array[$idx];
    $val
}
```

and

```
do {
    my $val = $hash{key};
    local $hash{key};
    delete $hash{key};
    $val
}
```

except that for those the `local` is scoped to the `do` block. Slices are also accepted.

```
my %hash = (
    a => [ 7, 8, 9 ],
    b => 1,
)

{
    my $a = delete local $hash{a};
    # $a is [ 7, 8, 9 ]
    # %hash is (b => 1)

    {
        my @nums = delete local @$a[0, 2]
        # @nums is (7, 9)
        # $a is [ undef, 8 ]

        $a[0] = 999; # will be erased when the scope ends
    }
    # $a is back to [ 7, 8, 9 ]

}
# %hash is back to its original state
```

Lvalue subroutines

It is possible to return a modifiable value from a subroutine. To do this, you have to declare the subroutine to return an lvalue.

```
my $val;
sub canmod : lvalue {
    $val; # or: return $val;
}
sub nomod {
    $val;
}

canmod() = 5; # assigns to $val
nomod()   = 5; # ERROR
```

The scalar/list context for the subroutine and for the right-hand side of assignment is determined as if the subroutine call is replaced by a scalar. For example, consider:

```
data(2,3) = get_data(3,4);
```

Both subroutines here are called in a scalar context, while in:

```
(data(2,3)) = get_data(3,4);
```

and in:

```
(data(2),data(3)) = get_data(3,4);
```

all the subroutines are called in a list context.

Lvalue subroutines are convenient, but you have to keep in mind that, when used with objects, they may violate encapsulation. A normal mutator can check the supplied argument before setting the

attribute it is protecting, an lvalue subroutine cannot. If you require any special processing when storing and retrieving the values, consider using the CPAN module Sentinel or something similar.

Lexical Subroutines

WARNING: Lexical subroutines are still experimental. The feature may be modified or removed in future versions of Perl.

Lexical subroutines are only available under the `use feature 'lexical_subs'` pragma, which produces a warning unless the "experimental::lexical_subs" warnings category is disabled.

Beginning with Perl 5.18, you can declare a private subroutine with `my` or `state`. As with state variables, the `state` keyword is only available under `use feature 'state'` or `use 5.010` or higher.

These subroutines are only visible within the block in which they are declared, and only after that declaration:

```
no warnings "experimental::lexical_subs";
use feature 'lexical_subs';

foo(); # calls the package/global subroutine
state sub foo {
foo(); # also calls the package subroutine
}
foo(); # calls "state" sub
my $ref = \&foo; # take a reference to "state" sub

my sub bar { ... }
bar(); # calls "my" sub
```

To use a lexical subroutine from inside the subroutine itself, you must predeclare it. The `sub foo { ... }` subroutine definition syntax respects any previous `my sub;` or `state sub;` declaration.

```
my sub baz; # predeclaration
sub baz { # define the "my" sub
baz(); # recursive call
}
```

state sub vs my sub

What is the difference between "state" subs and "my" subs? Each time that execution enters a block when "my" subs are declared, a new copy of each sub is created. "State" subroutines persist from one execution of the containing block to the next.

So, in general, "state" subroutines are faster. But "my" subs are necessary if you want to create closures:

```
no warnings "experimental::lexical_subs";
use feature 'lexical_subs';

sub whatever {
my $x = shift;
my sub inner {
... do something with $x ...
}
inner();
}
```


In this example, a new `$x` is created when `whatever` is called, and also a new `inner`, which can see the new `$x`. A "state" sub will only see the `$x` from the first call to `whatever`.

our subroutines

Like our `$variable`, our `sub` creates a lexical alias to the package subroutine of the same name.

The two main uses for this are to switch back to using the package sub inside an inner scope:

```
no warnings "experimental::lexical_subs";
use feature 'lexical_subs';

sub foo { ... }

sub bar {
my sub foo { ... }
{
    # need to use the outer foo here
    our sub foo;
    foo();
}
}
```

and to make a subroutine visible to other packages in the same scope:

```
package MySneakyModule;

no warnings "experimental::lexical_subs";
use feature 'lexical_subs';

our sub do_something { ... }

sub do_something_with_caller {
package DB;
() = caller 1; # sets @DB::args
do_something(@args); # uses MySneakyModule::do_something
}
```

Passing Symbol Table Entries (typeglobs)

WARNING: The mechanism described in this section was originally the only way to simulate pass-by-reference in older versions of Perl. While it still works fine in modern versions, the new reference mechanism is generally easier to work with. See below.

Sometimes you don't want to pass the value of an array to a subroutine but rather the name of it, so that the subroutine can modify the global copy of it rather than working with a local copy. In perl you can refer to all objects of a particular name by prefixing the name with a star: `*foo`. This is often known as a "typeglob", because the star on the front can be thought of as a wildcard match for all the funny prefix characters on variables and subroutines and such.

When evaluated, the typeglob produces a scalar value that represents all the objects of that name, including any filehandle, format, or subroutine. When assigned to, it causes the name mentioned to refer to whatever `*` value was assigned to it. Example:

```
sub doubleary {
local(*someary) = @_;
foreach $elem (@someary) {
```

```

    $elem *= 2;
}
}
doubleary(*foo);
doubleary(*bar);

```

Scalars are already passed by reference, so you can modify scalar arguments without using this mechanism by referring explicitly to `$_[0]` etc. You can modify all the elements of an array by passing all the elements as scalars, but you have to use the `*` mechanism (or the equivalent reference mechanism) to `push`, `pop`, or change the size of an array. It will certainly be faster to pass the typeglob (or reference).

Even if you don't want to modify an array, this mechanism is useful for passing multiple arrays in a single LIST, because normally the LIST mechanism will merge all the array values so that you can't extract out the individual arrays. For more on typeglobs, see *"Typeglobs and Filehandles" in perldata*.

When to Still Use local()

Despite the existence of `my`, there are still three places where the `local` operator still shines. In fact, in these three places, you *must* use `local` instead of `my`.

1. You need to give a global variable a temporary value, especially `$_`.

The global variables, like `@ARGV` or the punctuation variables, must be localized with `local()`. This block reads in */etc/motd*, and splits it up into chunks separated by lines of equal signs, which are placed in `@Fields`.

```

{
    local @ARGV = ("/etc/motd");
        local $/ = undef;
        local $_ = <>;
    @Fields = split /\s*+=+\s*$/;
}

```

It particular, it's important to localize `$_` in any routine that assigns to it. Look out for implicit assignments in `while` conditionals.

2. You need to create a local file or directory handle or a local function.

A function that needs a filehandle of its own must use `local()` on a complete typeglob. This can be used to create new symbol table entries:

```

sub ioqueue {
    local (*READER, *WRITER);    # not my!
    pipe (READER, WRITER)       or die "pipe: $!";
    return (*READER, *WRITER);
}
($head, $tail) = ioqueue();

```

See the Symbol module for a way to create anonymous symbol table entries.

Because assignment of a reference to a typeglob creates an alias, this can be used to create what is effectively a local function, or at least, a local alias.

```

{
    local *grow = \&shrink; # only until this block exits
    grow();                 # really calls shrink()
move();                    # if move() grow()'s, it shrink()'s too
}
grow();                    # get the real grow() again

```

See *"Function Templates" in perlref* for more about manipulating functions by name in this

- way.
- You want to temporarily change just one element of an array or hash.
- You can `localize` just one element of an aggregate. Usually this is done on dynamics:

```
{
local $SIG{INT} = 'IGNORE';
funct();          # uninterruptible
}
# interruptibility automatically restored here
```

But it also works on lexically declared aggregates.

Pass by Reference

If you want to pass more than one array or hash into a function--or return them from it--and have them maintain their integrity, then you're going to have to use an explicit pass-by-reference. Before you do that, you need to understand references as detailed in *perlref*. This section may not make much sense to you otherwise.

Here are a few simple examples. First, let's pass in several arrays to a function and have it `pop` all of them, returning a new list of all their former last elements:

```
@tailings = popmany ( \@a, \@b, \@c, \@d );

sub popmany {
my $aref;
my @retlist = ();
foreach $aref ( @_ ) {
    push @retlist, pop @$aref;
}
return @retlist;
}
```

Here's how you might write a function that returns a list of keys occurring in all the hashes passed to it:

```
@common = inter( \%foo, \%bar, \%joe );
sub inter {
my ($k, $href, %seen); # locals
foreach $href ( @_ ) {
    while ( $k = each %$href ) {
        $seen{$k}++;
    }
}
return grep { $seen{$_} == @_ } keys %seen;
}
```

So far, we're using just the normal list return mechanism. What happens if you want to pass or return a hash? Well, if you're using only one of them, or you don't mind them concatenating, then the normal calling convention is ok, although a little expensive.

Where people get into trouble is here:

```
(@a, @b) = func(@c, @d);
or
(%a, %b) = func(%c, %d);
```

That syntax simply won't work. It sets just `@a` or `%a` and clears the `@b` or `%b`. Plus the function didn't

get passed into two separate arrays or hashes: it got one long list in `@_`, as always.

If you can arrange for everyone to deal with this through references, it's cleaner code, although not so nice to look at. Here's a function that takes two array references as arguments, returning the two array elements in order of how many elements they have in them:

```
($aref, $bref) = func(\@c, \@d);
print "@$aref has more than @$bref\n";
sub func {
my ($cref, $dref) = @_;
if (@$cref > @$dref) {
    return ($cref, $dref);
} else {
    return ($dref, $cref);
}
}
```

It turns out that you can actually do this also:

```
(*a, *b) = func(\@c, \@d);
print "@a has more than @b\n";
sub func {
local (*c, *d) = @_;
if (@c > @d) {
    return (\@c, \@d);
} else {
    return (\@d, \@c);
}
}
```

Here we're using the typeglobs to do symbol table aliasing. It's a tad subtle, though, and also won't work if you're using `my` variables, because only globals (even in disguise as `locals`) are in the symbol table.

If you're passing around filehandles, you could usually just use the bare typeglob, like `*STDOUT`, but typeglobs references work, too. For example:

```
splutter(\*STDOUT);
sub splutter {
my $fh = shift;
print $fh "her um well a hmmm\n";
}

$rec = get_rec(\*STDIN);
sub get_rec {
my $fh = shift;
return scalar <$fh>;
}
```

If you're planning on generating new filehandles, you could do this. Notice to pass back just the bare `*FH`, not its reference.

```
sub openit {
my $path = shift;
local *FH;
return open (FH, $path) ? *FH : undef;
}
```

Prototypes

Perl supports a very limited kind of compile-time argument checking using function prototyping. This can be declared in either the PROTO section or with a *prototype attribute*. If you declare either of

```
sub mypush (+@)
sub mypush :prototype(+@)
```

then `mypush()` takes arguments exactly like `push()` does.

If subroutine signatures are enabled (see *Signatures*), then the shorter PROTO syntax is unavailable, because it would clash with signatures. In that case, a prototype can only be declared in the form of an attribute.

The function declaration must be visible at compile time. The prototype affects only interpretation of new-style calls to the function, where new-style is defined as not using the `&` character. In other words, if you call it like a built-in function, then it behaves like a built-in function. If you call it like an old-fashioned subroutine, then it behaves like an old-fashioned subroutine. It naturally falls out from this rule that prototypes have no influence on subroutine references like `&foo` or on indirect subroutine calls like `&{$subref}` or `$subref->()`.

Method calls are not influenced by prototypes either, because the function to be called is indeterminate at compile time, since the exact code called depends on inheritance.

Because the intent of this feature is primarily to let you define subroutines that work like built-in functions, here are prototypes for some other functions that parse almost exactly like the corresponding built-in.

Declared as	Called as
<code>sub mylink (\$\$)</code>	<code>mylink \$old, \$new</code>
<code>sub myvec (\$\$\$)</code>	<code>myvec \$var, \$offset, 1</code>
<code>sub myindex (\$\$;\$)</code>	<code>myindex &getstring, "substr"</code>
<code>sub mysyswrite (\$\$\$;\$)</code>	<code>mysyswrite \$buf, 0, length(\$buf) - \$off, \$off</code>
<code>sub myreverse (@)</code>	<code>myreverse \$a, \$b, \$c</code>
<code>sub myjoin (\$@)</code>	<code>myjoin ":", \$a, \$b, \$c</code>
<code>sub mypop (+)</code>	<code>mypop @array</code>
<code>sub mysplICE (+\$\$@)</code>	<code>mysplICE @array, 0, 2, @pushme</code>
<code>sub mykeys (+)</code>	<code>mykeys %{\$hashref}</code>
<code>sub myopen (*;\$)</code>	<code>myopen HANDLE, \$name</code>
<code>sub mypipe (**)</code>	<code>mypipe READHANDLE, WRITEHANDLE</code>
<code>sub mygrep (&@)</code>	<code>mygrep { /foo/ } \$a, \$b, \$c</code>
<code>sub myrand (;\$)</code>	<code>myrand 42</code>
<code>sub mytime ()</code>	<code>mytime</code>

Any backslashed prototype character represents an actual argument that must start with that character (optionally preceded by `my`, `our` or `local`), with the exception of `$`, which will accept any scalar lvalue expression, such as `$foo = 7` or `my_function()->[0]`. The value passed as part of `@_` will be a reference to the actual argument given in the subroutine call, obtained by applying `\` to that argument.

You can use the `\[]` backslash group notation to specify more than one allowed argument type. For example:

```
sub myref (\[$@%&*])
```

will allow calling `myref()` as

```
myref $var
myref @array
myref %hash
myref &sub
myref *glob
```

and the first argument of `myref()` will be a reference to a scalar, an array, a hash, a code, or a glob.

Unbackslashed prototype characters have special meanings. Any unbackslashed `@` or `%` eats all remaining arguments, and forces list context. An argument represented by `$` forces scalar context. An `&` requires an anonymous subroutine, which, if passed as the first argument, does not require the `sub` keyword or a subsequent comma.

A `*` allows the subroutine to accept a bareword, constant, scalar expression, typeglob, or a reference to a typeglob in that slot. The value will be available to the subroutine either as a simple scalar, or (in the latter two cases) as a reference to the typeglob. If you wish to always convert such arguments to a typeglob reference, use `Symbol::qualify_to_ref()` as follows:

```
use Symbol 'qualify_to_ref';

sub foo (*) {
my $fh = qualify_to_ref(shift, caller);
...
}
```

The `+` prototype is a special alternative to `$` that will act like `\[@%]` when given a literal array or hash variable, but will otherwise force scalar context on the argument. This is useful for functions which should accept either a literal array or an array reference as the argument:

```
sub mypush (+@) {
    my $aref = shift;
    die "Not an array or arrayref" unless ref $aref eq 'ARRAY';
    push @$aref, @_;
}
```

When using the `+` prototype, your function must check that the argument is of an acceptable type.

A semicolon (`;`) separates mandatory arguments from optional arguments. It is redundant before `@` or `%`, which gobble up everything else.

As the last character of a prototype, or just before a semicolon, a `@` or a `%`, you can use `_` in place of `$`: if this argument is not provided, `$_` will be used instead.

Note how the last three examples in the table above are treated specially by the parser. `mygrep()` is parsed as a true list operator, `myrand()` is parsed as a true unary operator with unary precedence the same as `rand()`, and `mytime()` is truly without arguments, just like `time()`. That is, if you say

```
mytime +2;
```

you'll get `mytime() + 2`, not `mytime(2)`, which is how it would be parsed without a prototype. If you want to force a unary function to have the same precedence as a list operator, add `;` to the end of the prototype:

```
sub mygetprotobynumber($;);
mygetprotobynumber $a > $b; # parsed as mygetprotobynumber($a > $b)
```

The interesting thing about `&` is that you can generate new syntax with it, provided it's in the initial

position:

```
    sub try (&@) {
my($try,$catch) = @_;
eval { &$try };
if ($@) {
    local $_ = $@;
    &$catch;
}
}
    sub catch (&) { $_[0] }

    try {
die "phooey";
    } catch {
/phooey/ and print "unphooey\n";
    };

```

That prints "unphooey". (Yes, there are still unresolved issues having to do with visibility of `@_`. I'm ignoring that question for the moment. (But note that if we make `@_` lexically scoped, those anonymous subroutines can act like closures... (Gee, is this sounding a little Lispish? (Never mind.))))

And here's a reimplementaion of the Perl `grep` operator:

```
    sub mygrep (&@) {
my $code = shift;
my @result;
foreach $_ (@_) {
    push(@result, $_) if &$code;
}
    @result;
    }

```

Some folks would prefer full alphanumeric prototypes. Alphanumerics have been intentionally left out of prototypes for the express purpose of someday in the future adding named, formal parameters. The current mechanism's main goal is to let module writers provide better diagnostics for module users. Larry feels the notation quite understandable to Perl programmers, and that it will not intrude greatly upon the meat of the module, nor make it harder to read. The line noise is visually encapsulated into a small pill that's easy to swallow.

If you try to use an alphanumeric sequence in a prototype you will generate an optional warning - "Illegal character in prototype...". Unfortunately earlier versions of Perl allowed the prototype to be used as long as its prefix was a valid prototype. The warning may be upgraded to a fatal error in a future version of Perl once the majority of offending code is fixed.

It's probably best to prototype new functions, not retrofit prototyping into older ones. That's because you must be especially careful about silent impositions of differing list versus scalar contexts. For example, if you decide that a function should take just one parameter, like this:

```
    sub func ($) {
my $n = shift;
print "you gave me $n\n";
    }

```

and someone has been calling it with an array or expression returning a list:

```
    func(@foo);

```

```
func( split /:/ );
```

Then you've just supplied an automatic `scalar` in front of their argument, which can be more than a bit surprising. The old `@foo` which used to hold one thing doesn't get passed in. Instead, `func()` now gets passed in a `1`; that is, the number of elements in `@foo`. And the `split` gets called in scalar context so it starts scribbling on your `@_` parameter list. Ouch!

If a sub has both a PROTO and a BLOCK, the prototype is not applied until after the BLOCK is completely defined. This means that a recursive function with a prototype has to be predeclared for the prototype to take effect, like so:

```
sub foo($$);
sub foo($$) {
    foo 1, 2;
}
```

This is all very powerful, of course, and should be used only in moderation to make the world a better place.

Constant Functions

Functions with a prototype of `()` are potential candidates for inlining. If the result after optimization and constant folding is either a constant or a lexically-scoped scalar which has no other references, then it will be used in place of function calls made without `&`. Calls made using `&` are never inlined. (See *constant.pm* for an easy way to declare most constants.)

The following functions would all be inlined:

```
sub pi () { 3.14159 } # Not exact, but close.
sub PI () { 4 * atan2 1, 1 } # As good as it gets,
    # and it's inlined, too!
sub ST_DEV () { 0 }
sub ST_INO () { 1 }

sub FLAG_FOO () { 1 << 8 }
sub FLAG_BAR () { 1 << 9 }
sub FLAG_MASK () { FLAG_FOO | FLAG_BAR }

sub OPT_BAZ () { not (0x1B58 & FLAG_MASK) }

sub N () { int(OPT_BAZ) / 3 }

sub FOO_SET () { 1 if FLAG_MASK & FLAG_FOO }
sub FOO_SET2 () { if (FLAG_MASK & FLAG_FOO) { 1 } }
```

(Be aware that the last example was not always inlined in Perl 5.20 and earlier, which did not behave consistently with subroutines containing inner scopes.) You can countermand inlining by using an explicit `return`:

```
sub baz_val () {
    if (OPT_BAZ) {
        return 23;
    }
    else {
        return 42;
    }
}
```



```
}  
sub bonk_val () { return 12345 }
```

As alluded to earlier you can also declare inlined subs dynamically at BEGIN time if their body consists of a lexically-scoped scalar which has no other references. Only the first example here will be inlined:

```
BEGIN {  
    my $var = 1;  
    no strict 'refs';  
    *INLINED = sub () { $var };  
}  
  
BEGIN {  
    my $var = 1;  
    my $ref = \"$var\";  
    no strict 'refs';  
    *NOT_INLINED = sub () { $var };  
}
```

A not so obvious caveat with this (see [RT #79908]) is that the variable will be immediately inlined, and will stop behaving like a normal lexical variable, e.g. this will print 79907, not 79908:

```
BEGIN {  
    my $x = 79907;  
    *RT_79908 = sub () { $x };  
    $x++;  
}  
print RT_79908(); # prints 79907
```

As of Perl 5.22, this buggy behavior, while preserved for backward compatibility, is detected and emits a deprecation warning. If you want the subroutine to be inlined (with no warning), make sure the variable is not used in a context where it could be modified aside from where it is declared.

```
# Fine, no warning  
BEGIN {  
    my $x = 54321;  
    *INLINED = sub () { $x };  
}  
# Warns. Future Perl versions will stop inlining it.  
BEGIN {  
    my $x;  
    $x = 54321;  
    *ALSO_INLINED = sub () { $x };  
}
```

Perl 5.22 also introduces the experimental "const" attribute as an alternative. (Disable the "experimental::const_attr" warnings if you want to use it.) When applied to an anonymous subroutine, it forces the sub to be called when the sub expression is evaluated. The return value is captured and turned into a constant subroutine:

```
my $x = 54321;  
*INLINED = sub : const { $x };  
$x++;
```

The return value of `INLINED` in this example will always be 54321, regardless of later modifications to `$x`. You can also put any arbitrary code inside the sub, at it will be executed immediately and its return value captured the same way.

If you really want a subroutine with a `()` prototype that returns a lexical variable you can easily force it to not be inlined by adding an explicit `return`:

```
BEGIN {
    my $x = 79907;
    *RT_79908 = sub () { return $x };
    $x++;
}
print RT_79908(); # prints 79908
```

The easiest way to tell if a subroutine was inlined is by using `B::Deparse`. Consider this example of two subroutines returning 1, one with a `()` prototype causing it to be inlined, and one without (with `deparse` output truncated for clarity):

```
$ perl -MO=Deparse -le 'sub ONE { 1 } if (ONE) { print ONE if ONE }'
sub ONE {
    1;
}
if (ONE ) {
    print ONE() if ONE ;
}
$ perl -MO=Deparse -le 'sub ONE () { 1 } if (ONE) { print ONE if ONE }'
sub ONE () { 1 }
do {
    print 1
};
```

If you redefine a subroutine that was eligible for inlining, you'll get a warning by default. You can use this warning to tell whether or not a particular subroutine is considered inlinable, since it's different than the warning for overriding non-inlined subroutines:

```
$ perl -e 'sub one () {1} sub one () {2}'
Constant subroutine one redefined at -e line 1.
$ perl -we 'sub one {1} sub one {2}'
Subroutine one redefined at -e line 1.
```

The warning is considered severe enough not to be affected by the `-w` switch (or its absence) because previously compiled invocations of the function will still be using the old value of the function. If you need to be able to redefine the subroutine, you need to ensure that it isn't inlined, either by dropping the `()` prototype (which changes calling semantics, so beware) or by thwarting the inlining mechanism in some other way, e.g. by adding an explicit `return`, as mentioned above:

```
sub not_inlined () { return 23 }
```

Overriding Built-in Functions

Many built-in functions may be overridden, though this should be tried only occasionally and for good reason. Typically this might be done by a package attempting to emulate missing built-in functionality on a non-Unix system.

Overriding may be done only by importing the name from a module at compile time--ordinary predeclaration isn't good enough. However, the `use subs` pragma lets you, in effect, predeclare subs via the import syntax, and these names may then override built-in ones:

```
use subs 'chdir', 'chroot', 'chmod', 'chown';
chdir $somewhere;
sub chdir { ... }
```

To unambiguously refer to the built-in form, precede the built-in name with the special package qualifier `CORE::`. For example, saying `CORE::open()` always refers to the built-in `open()`, even if the current package has imported some other subroutine called `&open()` from elsewhere. Even though it looks like a regular function call, it isn't: the `CORE::` prefix in that case is part of Perl's syntax, and works for any keyword, regardless of what is in the `CORE` package. Taking a reference to it, that is, `\&CORE::open`, only works for some keywords. See *CORE*.

Library modules should not in general export built-in names like `open` or `chdir` as part of their default `@EXPORT` list, because these may sneak into someone else's namespace and change the semantics unexpectedly. Instead, if the module adds that name to `@EXPORT_OK`, then it's possible for a user to import the name explicitly, but not implicitly. That is, they could say

```
use Module 'open';
```

and it would import the `open` override. But if they said

```
use Module;
```

they would get the default imports without overrides.

The foregoing mechanism for overriding built-in is restricted, quite deliberately, to the package that requests the import. There is a second method that is sometimes applicable when you wish to override a built-in everywhere, without regard to namespace boundaries. This is achieved by importing a sub into the special namespace `CORE::GLOBAL::`. Here is an example that quite brazenly replaces the `glob` operator with something that understands regular expressions.

```
package REGlob;
require Exporter;
@ISA = 'Exporter';
@EXPORT_OK = 'glob';

sub import {
    my $pkg = shift;
    return unless @_;
    my $sym = shift;
    my $where = ($sym =~ s/^GLOBAL_// ? 'CORE::GLOBAL' : caller(0));
    $pkg->export($where, $sym, @_);
}

sub glob {
    my $pat = shift;
    my @got;
    if (opendir my $d, '.') {
        @got = grep /$pat/, readdir $d;
        closedir $d;
    }
    return @got;
}
1;
```

And here's how it could be (ab)used:

```
#use REGlob 'GLOBAL_glob';      # override glob() in ALL namespaces
package Foo;
use REGlob 'glob';              # override glob() in Foo:: only
print for <^[a-z_]+\.\pm\$>;      # show all pragmatic modules
```

The initial comment shows a contrived, even dangerous example. By overriding `glob` globally, you would be forcing the new (and subversive) behavior for the `glob` operator for every namespace, without the complete cognizance or cooperation of the modules that own those namespaces. Naturally, this should be done with extreme caution--if it must be done at all.

The `REGlob` example above does not implement all the support needed to cleanly override perl's `glob` operator. The built-in `glob` has different behaviors depending on whether it appears in a scalar or list context, but our `REGlob` doesn't. Indeed, many perl built-in have such context sensitive behaviors, and these must be adequately supported by a properly written override. For a fully functional example of overriding `glob`, study the implementation of `File::DosGlob` in the standard library.

When you override a built-in, your replacement should be consistent (if possible) with the built-in native syntax. You can achieve this by using a suitable prototype. To get the prototype of an overridable built-in, use the `prototype` function with an argument of `"CORE::builtin_name"` (see *"prototype" in perlfunc*).

Note however that some built-ins can't have their syntax expressed by a prototype (such as `system` or `chomp`). If you override them you won't be able to fully mimic their original syntax.

The built-ins `do`, `require` and `glob` can also be overridden, but due to special magic, their original syntax is preserved, and you don't have to define a prototype for their replacements. (You can't override the `do BLOCK` syntax, though).

`require` has special additional dark magic: if you invoke your `require` replacement as `require Foo::Bar`, it will actually receive the argument `"Foo/Bar.pm"` in `@_`. See *"require" in perlfunc*.

And, as you'll have noticed from the previous example, if you override `glob`, the `<*>` `glob` operator is overridden as well.

In a similar fashion, overriding the `readline` function also overrides the equivalent I/O operator `<FILEHANDLE>`. Also, overriding `readpipe` also overrides the operators `` `` and `qx//`.

Finally, some built-ins (e.g. `exists` or `grep`) can't be overridden.

Autoloading

If you call a subroutine that is undefined, you would ordinarily get an immediate, fatal error complaining that the subroutine doesn't exist. (Likewise for subroutines being used as methods, when the method doesn't exist in any base class of the class's package.) However, if an `AUTOLOAD` subroutine is defined in the package or packages used to locate the original subroutine, then that `AUTOLOAD` subroutine is called with the arguments that would have been passed to the original subroutine. The fully qualified name of the original subroutine magically appears in the global `$AUTOLOAD` variable of the same package as the `AUTOLOAD` routine. The name is not passed as an ordinary argument because, er, well, just because, that's why. (As an exception, a method call to a nonexistent `import` or `unimport` method is just skipped instead. Also, if the `AUTOLOAD` subroutine is an `XSUB`, there are other ways to retrieve the subroutine name. See *"Autoloading with XSUBs" in perl guts* for details.)

Many `AUTOLOAD` routines load in a definition for the requested subroutine using `eval()`, then execute that subroutine using a special form of `goto()` that erases the stack frame of the `AUTOLOAD` routine without a trace. (See the source to the standard module documented in *AutoLoader*, for example.) But an `AUTOLOAD` routine can also just emulate the routine and never define it. For example, let's pretend that a function that wasn't defined should just invoke `system` with those arguments. All you'd do is:

```
sub AUTOLOAD {
my $program = $AUTOLOAD;
$program =~ s/.*:://;
system($program, @_);
}
date();
who('am', 'i');
ls('-l');
```

In fact, if you predeclare functions you want to call that way, you don't even need parentheses:

```
use subs qw(date who ls);
date;
who "am", "i";
ls '-l';
```

A more complete example of this is the *Shell* module on CPAN, which can treat undefined subroutine calls as calls to external programs.

Mechanisms are available to help modules writers split their modules into autoloadable files. See the standard *AutoLoader* module described in *AutoLoader* and in *AutoSplit*, the standard *SelfLoader* modules in *SelfLoader*, and the document on adding C functions to Perl code in *perlxs*.

Subroutine Attributes

A subroutine declaration or definition may have a list of attributes associated with it. If such an attribute list is present, it is broken up at space or colon boundaries and treated as though a `use attributes` had been seen. See *attributes* for details about what attributes are currently supported. Unlike the limitation with the obsolescent `use attrs`, the `sub : ATTRLIST` syntax works to associate the attributes with a pre-declaration, and not just with a subroutine definition.

The attributes must be valid as simple identifier names (without any punctuation other than the `'_'` character). They may have a parameter list appended, which is only checked for whether its parentheses `('(',')')` nest properly.

Examples of valid syntax (even though the attributes are unknown):

```
sub fnord (&\%) : switch(10,foo(7,3)) : expensive;
sub plugh () : Ugly('\(") :Bad;
sub xyzzy : _5x5 { ... }
```

Examples of invalid syntax:

```
sub fnord : switch(10,foo()); # ()-string not balanced
sub snoid : Ugly('('); # ()-string not balanced
sub xyzzy : 5x5; # "5x5" not a valid identifier
sub plugh : Y2::north; # "Y2::north" not a simple identifier
sub snurt : foo + bar; # "+" not a colon or space
```

The attribute list is passed as a list of constant strings to the code which associates them with the subroutine. In particular, the second example of valid syntax above currently looks like this in terms of how it's parsed and invoked:

```
use attributes __PACKAGE__, \&plugh, q[Ugly('\(")], 'Bad';
```

For further details on attribute lists and their manipulation, see *attributes* and *Attribute::Handlers*.

SEE ALSO

See "*Function Templates*" in *perlref* for more about references and closures. See *perlxs* if you'd like to learn about calling C subroutines from Perl. See *perlembed* if you'd like to learn about calling Perl subroutines from C. See *perlmod* to learn about bundling up your functions in separate files. See *perlmodlib* to learn what library modules come standard on your system. See *perloutut* to learn how to make object method calls.