

The GNU Shepherd Manual

For use with the GNU Shepherd 0.9.2
Last updated 10 September 2022

Wolfgang Jährling
Ludovic Courtès

Copyright © 2002, 2003 Wolfgang Jährling

Copyright © 2013, 2016, 2018, 2019, 2020, 2022 Ludovic Courtès

Copyright © 2020 Brice Waegeneire

Copyright © 2020 Oleg Pykhalov Copyright © 2020 Jan (janneke) Nieuwenhuizen

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.3 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled “GNU Free Documentation License”.

Table of Contents

1	Introduction	1
2	Jump Start	2
3	herd and shepherd	5
3.1	Invoking <code>shepherd</code>	5
3.2	Invoking <code>herd</code>	6
3.3	Invoking <code>reboot</code>	7
3.4	Invoking <code>halt</code>	7
4	Services	8
4.1	Slots of services	8
4.2	Methods of services	10
4.3	Service Convenience	11
4.4	Service De- and Constructors	11
4.5	Service Examples	15
4.6	Managing User Services	17
4.7	The <code>root</code> and <code>unknown</code> services	18
5	Misc Facilities	20
5.1	Errors	20
5.2	Communication	20
5.3	Others	21
6	Internals	22
6.1	Coding standards	22
6.2	Design decisions	22
6.3	Service Internals	24
	Appendix A GNU Free Documentation License ..	25
	Concept Index	33
	Procedure and Macro Index	35
	Variable Index	36
	Type Index	37

1 Introduction

This manual documents the GNU Daemon Shepherd, or GNU Shepherd for short. The Shepherd looks after system services, typically *daemons*. It is used to start and stop them in a reliable fashion. For instance it will dynamically determine and start any other services that our desired service depends upon. As another example, the Shepherd might detect conflicts among services. In this situation it would simply prevent the conflicting services from running concurrently.

The Shepherd is the *init system* of the GNU operating system—it is the first user process that gets started, typically with PID 1, and runs as `root`. Normally the purpose of *init* systems is to manage all system-wide services, but the Shepherd can also be a useful tool assisting unprivileged users in the management of their own daemons.

Flexible software requires some time to master and the Shepherd is no different. But don't worry: this manual should allow you to get started quickly. Its first chapter is designed as a practical introduction to the Shepherd and should be all you need for everyday use (see Chapter 2 [Jump Start], page 2). In chapter two we will describe the `herd` and `shepherd` programs, and their relationship, in more detail (Chapter 3 [herd and shepherd], page 5). Subsequent chapters provide a full reference manual and plenty of examples, covering all of Shepherd's capabilities. Finally, the last chapter provides information for those souls brave enough to hack the Shepherd itself.

The Shepherd was formerly known as “`dmd`”, which stands for *Daemon Managing Daemons* (or *Daemons-Managing Daemon?*).

This program is written in Guile, an implementation of the Scheme programming language, using the GOOPS extension for object-orientation. Guile is also the Shepherd's configuration language. See Section “Introduction” in *GNU Guile Reference Manual*, for an introduction to Guile. We have tried to make the Shepherd's basic features as accessible as possible—you should be able to use these even if you do not know how to program in Scheme. A basic grasp of Guile and GOOPS is required only if you wish to make use of the Shepherd's more advanced features.

2 Jump Start

This chapter gives a short overview of the Shepherd. It is enough if you just need the basic features of it. As it is not assumed that readers are familiar with all the involved issues, a very experienced user might be annoyed by the often very detailed descriptions in this introduction. Those users are encouraged to just skip to the reference section.

Note that all the full file names in the following text are based on the assumption that you have installed the Shepherd with an empty prefix. If your Shepherd installation for example resides in `/usr/local` instead, add this directory name in front of the absolute file names mentioned below.

When `shepherd` gets started, it reads and evaluates a configuration file. When it is started with superuser privileges, it tries to use `/etc/shepherd.scm`. When started as normal user, it looks for a file called `$XDG_CONFIG_HOME/shepherd/init.scm`. If the `XDG_CONFIG_HOME` environment variable is not defined, `$HOME/.config/shepherd/init.scm` is used instead (see [\[Managing User Services\]](#), page [\(undefined\)](#)). With the option `--config` (or, for short, `-c`), you can specify where to look instead. So if you want to start `shepherd` with an alternative file, use one of the following commands:

```
shepherd --config=/etc/shepherd.scm.old
shepherd -c /etc/shepherd.scm.old
```

As the final “d” suggests, `shepherd` is just a daemon that (usually) runs in the background, so you will not interact with it directly. After it is started, `shepherd` will listen on a socket special file, usually `/var/run/shepherd/socket`, for further commands. You use the tool `herd` to send these commands to `shepherd`. Usage of `herd` is simple and straightforward: To start a service called `apache`, you use:

```
herd start apache
```

When you do this, all its dependencies will get resolved. For example, a webserver is quite likely to depend on working networking, thus it will depend on a service called `networking`. So if you want to start `apache`, and `networking` is not yet running, it will automatically be started as well. The current status of all the services defined in the configuration file can be queried like this:

```
herd status
```

Or, to get additional details about each service, run:

```
herd detailed-status
```

In this example, this would show the `networking` and `apache` services as started. If you just want to know the status of the `apache` service, run:

```
herd status apache
```

You can stop a service and all the services that depend on it will be stopped. Using the example above, if you stop `networking`, the service `apache` will be stopped as well—which makes perfect sense, as it cannot work without the network being up. To actually stop a service, you use the following, probably not very surprising, command:

```
herd stop networking
```

There are two more actions you can perform on every service: the actions `enable` and `disable` are used to prevent and allow starting of the particular service. If a service is intended to be restarted whenever it terminates (how this can be done will not be covered

in this introduction), but it is respawning too often in a short period of time (by default 5 times in 5 seconds), it will automatically be disabled. After you have fixed the problem that caused it from being respawned too fast, you can start it again with the commands:

```
herd enable foo
herd start foo
```

But there is far more you can do than just that. Services can not only simply depend on other services, they can also depend on *virtual* services. A virtual service is a service that is provided by one or more service additionally. For instance, a service called `exim` might provide the virtual service `mailer-daemon`. That could as well be provided by a service called `smail`, as both are mailer-daemons. If a service needs any mailer-daemon, no matter which one, it can just depend on `mailer-daemon`, and one of those who provide it gets started (if none is running yet) when resolving dependencies. The nice thing is that, if trying to start one of them fails, `shepherd` will go on and try to start the next one, so you can also use virtual services for specifying *fallbacks*.

Additionally to all that, you can perform service-specific actions. Coming back to our original example, `apache` is able to reload its modules, therefore the action `reload-modules` might be available:

```
herd reload-modules apache
```

Service-specific actions can only be used when the service is started, i.e. the only thing you can do to a stopped service is starting it. An exception exists, see below. (If you may at some point find this too restrictive because you want to use variants of the same service which are started in different ways, consider using different services for those variants instead, which all provide the same virtual service and thus conflict with each other, if this is desired. That's one of the reasons why virtual services exist, after all.)

There are two actions which are special, because even if services can implement them on their own, a default implementation is provided by `shepherd` (another reason why they are special is that the default implementations can be called even when the service is not running; this inconsistency is just to make it more intuitive to get information about the status of a service, see below).

These actions are `restart` and `status`. The default implementation of `restart` calls `stop` and `start` on the affected service, taking care to also restart any dependent services. The default implementation of `status` displays some general information about the service, like what it provides, what it depends on and with which other services it conflicts (because they provide a virtual service that is also provided by that particular service).

A special service is `root`, which is used for controlling the Shepherd itself. You can also reference to this service as `shepherd`. It implements various actions. For example, the `status` action displays which services are started and which ones are stopped, whereas `detailed-status` has the effect of applying the default implementation of `status` to all services one after another. The `load` action is unusual insofar as it shows a feature that is actually available to all services, but which we have not seen yet: It takes an additional argument. You can use `load` to load arbitrary code into the Shepherd at runtime, like this:

```
herd load shepherd ~/additional-services.scm
```

In the same vein the special action `doc` describes its service when called without an argument or describes a service-specific action when called with the action as the additional

arguments. You can even get the list of the service-specific actions a service provides when using with the additional argument `list-actions`.

```
$ herd doc root
```

```
The root service is used to operate on shepherd itself.
```

```
$ herd doc root list-actions
```

```
root (help status halt power-off load eval unload reload daemonize persistency no-pers
```

```
$ herd doc root action power-off
```

```
power-off: Halt the system and turn it off.
```

This is enough now about the `herd` and `shepherd` programs, we will now take a look at how to configure the Shepherd. In the configuration file, we need mainly the definition of services. We can also do various other things there, like starting a few services already.

FIXME: Finish. For now, look at the `doc/examples/` subdirectory.

```
...
```

Ok, to summarize:

- `shepherd` is a daemon, `herd` the program that controls it.
- You can start, stop, restart, enable and disable every service, as well as display its status.
- You can perform additional service-specific actions, which you can also list.
- Actions can have arguments.
- You can display the status of a service, even if the service does not provide a specific implementation for this action. The same is true for restarting.
- The `root/shepherd` service is used to control `shepherd` itself.

3 herd and shepherd

The daemon that runs in the background and is responsible for controlling the services is **shepherd**, while the user interface tool is called **herd**: it's the command that allows you to actually *herd* your daemons¹. To perform an action, like stopping a service or calling an action of a service, you use the herd program. It will communicate with shepherd over a Unix Domain Socket.

Thus, you start **shepherd** once, and then always use herd whenever you want to do something service-related. Since herd passes its current working directory to **shepherd**, you can pass relative file names without trouble. Both **shepherd** and herd understand the standard arguments `--help`, `--version` and `--usage`.

3.1 Invoking shepherd

The **shepherd** program has the following synopsis:

```
shepherd [option...]
```

It accepts the following options:

```
'-c file'
```

```
'--config=file'
```

Read and evaluate *file* as the configuration script on startup.

file is evaluated in the context of a fresh module where bindings from the (**shepherd service**) module and Guile's (`oop goops`) are available, in addition to the default set of Guile bindings. In particular, this means that code in *file* may use `register-services`, the `<service>` class, and related tools (see Chapter 4 [Services], page 8).

```
'-I'
```

```
'--insecure'
```

Do not check if the directory where the socket—our communication rendezvous with **herd**—is located has permissions 700. If this option is not specified, **shepherd** will abort if the permissions are not as expected.

```
'-l [file]'
```

```
'--logfile[=file]'
```

Log output into *file*.

For unprivileged users, the default log file is `$XDG_STATE_HOME/shepherd/shepherd.log` with `$XDG_STATE_HOME` defaulting to `$HOME/.local/state`.

When running as root, the default behavior is to connect to `/dev/log`, the *syslog* socket (see Section “Overview of Syslog” in *The GNU C Library Reference Manual*). A *syslog* daemon, `syslogd`, is expected to read messages from there (see Section “`syslogd` invocation” in *GNU Inetutils*).

When `/dev/log` is unavailable, for instance because `syslogd` is not running, as is the case during system startup and shutdown, **shepherd** falls back to the

¹ In the past, when the GNU Shepherd was known as GNU dmd, the **herd** command was called **deco**, for *DaEmon COntroller*.

Linux kernel *ring buffer*, `/dev/kmsg`. If `/dev/kmsg` is missing, as is the case on other operating systems, it falls back to `/dev/console`.

`--pid[=file]`

When `shepherd` is ready to accept connections, write its PID to *file* or to the standard output if *file* is omitted.

`-p [file]`

`--persistency[=file]`

`-s file`

`--socket=file`

Receive further commands on the socket special file *file*. If this option is not specified, `localstatedir/run/shepherd/socket` is taken when running as `root`; when running as an unprivileged user, `shepherd` listens to `/run/user/uid/shepherd/socket`, where *uid* is the user's numerical ID², or to `$XDG_RUNTIME_DIR/shepherd` when the `XDG_RUNTIME_DIR` environment variable is defined.

If `-` is specified as file name, commands will be read from standard input, one per line, as would be passed on a `herd` command line (see Section 3.2 [Invoking herd], page 6).

`--quiet` Synonym for `--silent`.

3.2 Invoking herd

The `herd` command is a generic client program to control a running instance of `shepherd` (see Section 3.1 [Invoking shepherd], page 5). When running as `root`, it communicates with the *system instance*—the process with PID 1; when running as a normal user, it communicates with the *user's instance*, which is a regular, unprivileged process managing the user's own services. For example, the following command displays the status of all the *system services*:

```
sudo herd status
```

Conversely, the command below displays the status of *user services*, assuming a user `shepherd` is running:

```
herd status
```

The command has the following synopsis:

```
herd [option...] action [service [arg...]]
```

It causes the *action* of the *service* to be invoked. When *service* is omitted and *action* is `status` or `detailed-status`, the `root` service is used³ (see Section 4.7 [The root and unknown services], page 18, for more information on the `root` service.)

For each action, you should pass the appropriate *args*. Actions that are available for every service are `start`, `stop`, `restart`, `status`, `enable`, `disable`, and `doc`.

² On GNU/Linux, the `/run/user/uid` directory is typically created by `elogind` or by `systemd`, which are available in most distributions.

³ This shorthand does not work for other actions such as `stop`, because inadvertently typing `herd stop` would stop all the services, which could be pretty annoying.

If you pass a file name as an *arg*, it will be passed as-is to the Shepherd, thus if it is not an absolute name, it is local to the current working directory of **shepherd**, not to **herd**.

The **herd** command understands the following option:

`'-s file'`

`'--socket=file'`

Send commands to the socket special file *file*. If this option is not specified, `localstatedir/run/shepherd/socket` is taken.

The **herd** command returns zero on success, and a non-zero exit code on failure. In particular, it returns a non-zero exit code when *action* or *service* does not exist and when the given action failed.

3.3 Invoking reboot

The **reboot** command is a convenience client program to instruct the Shepherd (when used as an init system) to stop all running services and reboot the system. It has the following synopsis:

```
reboot [option...]
```

It is equivalent to running `herd stop shepherd`. The **reboot** command understands the following option:

`'-s file'`

`'--socket=file'`

Send commands to the socket special file *file*. If this option is not specified, `localstatedir/run/shepherd/socket` is taken.

3.4 Invoking halt

The **halt** command is a convenience client program to instruct the Shepherd (when used as an init system) to stop all running services and turn off the system. It has the following synopsis:

```
halt [option...]
```

It is equivalent to running `herd power-off shepherd`. As usual, the **halt** command understands the following option:

`'-s file'`

`'--socket=file'`

Send commands to the socket special file *file*. If this option is not specified, `localstatedir/run/shepherd/socket` is taken.

4 Services

The *service* is obviously a very important concept of the Shepherd. On the Guile level, a service is represented as an instance of `<service>`, a GOOPS class (see Section “GOOPS” in *GNU Guile Reference Manual*). When creating an instance of it, you can specify the initial values of its slots, and you actually must do this for some of the slots.

The `<service>` class and its associated procedures and methods are defined in the `(shepherd service)` module.

4.1 Slots of services

A service has the following slots, all of which can be initialized with a keyword (i.e. `#:provides`, used when creating the object) of the same name, except where stated otherwise. You should not access them directly with `slot-ref` or `slot-set!` usually, use the methods of the service class Section 4.2 [Methods of services], page 10, instead.

- `provides` is a list of symbols that are provided by the service. A symbol can only be provided by one service running at a time, i.e. if two services provide the same symbol, only one of them can run, starting the other one will fail. Therefore, these symbols are mainly used to denote conflicting services. The first symbol in the list is the canonical name for the service, thus it must be unique. This slot has no default value and must therefore be initialized.
- `requires` is, like `provides`, a list of symbols that specify services. In this case, they name what this service depends on, i.e. before the service can be started, services that provide those symbols must be started. If a required symbol is provided by several services, one will be started. By default, this slot contains the empty list.
- `running` is a hook that can be used by each service in its own way. The default value is `#f`, which indicates that the service is not running. When an attempt is made to start the service, it will be set to the return value of the procedure in the `start` slot. It will also be passed as an argument to the procedure in the `stop` slot. If it is set a value that is an integer, it is assumed to be a process id, and shepherd will monitor the process for unexpected exits. If it is a procedure, that procedure is called to get at the underlying value. This slot cannot be initialized with a keyword.
- `respawn?` specifies whether the service should be respawned by the Shepherd. If this slot has the value `#t`, then assume the `running` slot specifies a child process PID and restart the service if that process terminates. Otherwise this slot is `#f`, which is the default. See also the `last-respawns` slot.
- The `one-shot?` slot determines whether the service is a *one-shot service*. A one-shot service is a service that, as soon as it has been successfully started, is marked as “stopped.” Other services can nonetheless require one-shot services. One-shot services are useful to trigger an action before other services are started, such as a cleanup or an initialization action.

As for other services, the `start` method of a one-shot service must return a truth value to indicate success, and false to indicate failure.

- The `transient?` slot determines whether the service is a *transient service*. A transient service is automatically unregistered when it terminates, be it because its `stop` method is called or because its associated process terminates.

This is useful in the uncommon case of synthesized services that may not be restarted once they have completed.

- **start** contains the “constructor” for the service, which will be called to start the service. (Therefore, it is not a constructor in the sense that it initializes the slots of a `<service>` object.) This must be a procedure that accepts any amount of arguments, which will be the additional arguments supplied by the user. If the starting attempt failed, it must return `#f`. The value will be stored in the **running** slot. The default value is a procedure that returns `#t` and performs no further actions, therefore it is desirable to specify a different one usually.
- **stop** is, similar to **start**, a slot containing a procedure. But in this case, it gets the current value of the **running** slot as first argument and the user-supplied arguments as further arguments; it gets called to stop the service. Its return value will again be stored in the **running** slot, so that it should return `#f` if it is now possible again to start the service at a later point. The default value is a procedure that returns `#f` and performs no further actions.
- The **handle-termination** slot contains the procedure to call when the process associated with a service—the process whose PID appears in the **running** slot—terminates. It is passed the service and its exit status, an integer as returned by `waitpid` (see Section “Processes” in *GNU Guile Reference Manual*).

The default handler is the `default-service-termination-handler` procedure, which respawns the service if applicable.

- **actions** specifies the additional actions that can be performed on a service when it is running. A typical example for this is the **restart** action. The macro `make-actions` Section 4.3 [Service Convenience], page 11, is provided to abstract the actual data representation format for this slot. (It actually is a hash currently.)
- **enabled?** cannot be initialized with a keyword, and contains `#t` by default. When the value becomes `#f` at some point, this will prevent the service from getting started. A service can be enabled and disabled with the methods `enable` and `disable`, respectively Section 4.2 [Methods of services], page 10.
- **last-respawns** cannot be initialized with a keyword and is only ever used when the **respawn?** slot contains `#t`; it is a circular list with `(car respawn-limit)` elements, where each element contains the time when it was restarted, initially all 0, later a time in seconds since the Epoch. The first element is the one that contains the oldest one, the last one the newest.
- **stop-delay?** being false causes the **stop** slot to be unused; instead, stopping the service will just cause the **waiting-for-termination?** slot be set to `#t`.
- **waiting-for-termination?** cannot be initialized with a keyword and should not be used by others, it is only used internally for respawnable services when the **stop-delay?** slot contains a true value. **waiting-for-termination?** contains `#t` if the service is still running, but the user requested that it be stopped, in which case if the service terminates the next time, the respawn handler will not start it again.
otherwise `#f`.
- **replacement** specifies a service to be used to replace this one when it is stopped. This service will continue to function normally until the **stop** action is invoked. After the

service has been successfully stopped, its definition will be replaced by the value of this slot, which must itself be a service. This slot is ignored if its value is `#f`.

4.2 Methods of services

`start (obj <service>)` [method]

Start the service *obj*, including all the services it depends on. It tries quite hard to do this: When a service that provides a required symbol can not be started, it will look for another service that also provides this symbol, until starting one such service succeeds. There is some room for theoretical improvement here, of course, but in practice the current strategy already works very well. This method returns the new value of the `running` slot Section 4.1 [Slots of services], page 8, which is `#f` if the service could not be started.

`stop (obj <service>)` [method]

This will stop the service *obj*, trying to stop services that depend in it first, so they can be shutdown cleanly. If this will fail, it will continue anyway. Stopping of services should usually succeed, though. Otherwise, the behaviour is very similar to the `start` method. The return value is also the new `running` value, thus `#f` if the service was stopped.

`action (obj <service>) the-action . args` [method]

Calls the action *the-action* (a symbol) of the service *obj*, with the specified *args*, which have a meaning depending on the particular action.

`conflicts-with (obj <service>)` [method]

Returns a list of the canonical names of services that conflict with the service *obj*.

`canonical-name (obj <service>)` [method]

Returns the canonical name of *obj*, which is the first element of the `provides` list.

`provided-by (obj <service>)` [method]

Returns which symbols are provided by *obj*.

`required-by (obj <service>)` [method]

Returns which symbols are required by *obj*.

`one-shot? (obj <service>)` [method]

Returns whether the service *obj* is a one-shot service.

`running? (obj <service>)` [method]

Returns whether the service *obj* is running.

`respawn? (obj <service>)` [method]

Returns whether the service *obj* should be respawned if it terminates.

`default-display-status (obj <service>)` [method]

Display status information about *obj*. This method is called when the user performs the action `status` on *obj*, but there is no specific implementation given for it. It is also called when `detailed-status` is applied on the `root` service.

4.3 Service Convenience

In addition to the facilities listed below, there are also some procedures that provide commonly needed constructors and destructors for services Section 4.4 [Service De- and Constructors], page 11.

`register-services . services` [procedure]
 Register all *services*, so that they can be taken into account when trying to resolve dependencies.

`lookup-services name` [procedure]
 Return a list of all registered services which provide the symbol *name*.

`make-actions (name proc) ...` [macro]
 This macro is used to create a value for the `actions` slot of a service object Section 4.1 [Slots of services], page 8. Each *name* is a symbol and each *proc* the corresponding procedure that will be called to perform the action. A *proc* has one argument, which will be the current value of the `running` slot of the service.

`start (obj <symbol>)` [method]
 Start a registered service providing *obj*.

`start-in-the-background services` [procedure]
 Start the services named by *services*, a list of symbols, in the background. In other words, this procedure returns immediately without waiting until all of *services* have been started.

This procedure can be useful in a configuration file because it lets you interact right away with shepherd using the `herd` command.

`stop (obj <symbol>)` [method]
 Stop a registered service providing *obj*.

`action (obj <symbol>) the-action . args` [method]
 The same as the `action` method of class `<service>`, but uses a service that provides *obj* and is running.

`for-each-service proc` [procedure]
 Call *proc*, a procedure taking one argument, once for each registered service.

`find-running services` [procedure]
 Check if any of *services* is running. If this is the case, return its canonical name. If not, return `#f`. Only the first one will be returned; this is because this is mainly intended to be applied on the return value of `lookup-services`.

4.4 Service De- and Constructors

All of the procedures listed below return procedures generated from the supplied arguments. These procedures take one argument in the case of destructors and no arguments in the case of constructors.

`make-system-constructor` *command*... [procedure]

The returned procedure will execute *command* in a shell and return `#t` if execution was successful, otherwise `#f`. For convenience, it takes multiple arguments which will be concatenated first.

`make-system-destructor` *command*... [procedure]

Similar to `make-system-constructor`, but returns `#f` if execution of the *command* was successful, `#t` if not.

`make-forkexec-creator` *command* [`#:user #f`] [`#:group #f`] [procedure]
 [`#:supplementary-groups '()`] [`#:pid-file #f`] [`#:pid-file-timeout`
 (*default-pid-file-timeout*)] [`#:log-file #f`] [`#:directory`
 (*default-service-directory*)] [`#:file-creation-mask #f`] [`#:create-session?`
`#t`] [`#:resource-limits '()`] [`#:environment-variables`
 (*default-environment-variables*)]

Return a procedure that forks a child process, closes all file descriptors except the standard output and standard error descriptors, sets the current directory to *directory*, sets the umask to *file-creation-mask* unless it is `#f`, changes the environment to *environment-variables* (using the `environ` procedure), sets the current user to *user* the current group to *group* unless they are `#f` and supplementary groups to *supplementary-groups* unless they are `'()`, and executes *command* (a list of strings.) When *create-session?* is true, the child process creates a new session with `setsid` and becomes its leader. The result of the procedure will be the PID of the child process.

Note: This will not work as expected if the process “daemonizes” (forks); in that case, you will need to pass `#:pid-file`, as explained below.

When *pid-file* is true, it must be the name of a PID file associated with the process being launched; the return value is the PID once that file has been created. If *pid-file* does not show up in less than *pid-file-timeout* seconds, the service is considered as failing to start.

When *log-file* is true, it names the file to which the service’s standard output and standard error are redirected. *log-file* is created if it does not exist, otherwise it is appended to.

Guile’s `setrlimit` procedure is applied on the entries in *resource-limits*. For example, a valid value would be:

```
'((nproc 10 100)           ;number of processes
  (nofile 4096 4096))     ;number of open file descriptors
```

`make-kill-destructor` [*signal*] [procedure]

Return a procedure that sends *signal* to the process group of the PID given as argument, where *signal* defaults to `SIGTERM`.

This *does* work together with respawning services, because in that case the `stop` method of the `<service>` class sets the `running` slot to `#f` before actually calling the destructor; if it would not do that, killing the process in the destructor would immediately respawn the service.

The `make-forkexec-creator` procedure builds upon the following procedures.

```

exec-command command [#:user #f] [#:group #f] [procedure]
  [#:supplementary-groups '()] [#:log-file #f] [#:log-port #f]
  [#:input-port #f] [#:directory (default-service-directory)]
  [#:file-creation-mask #f] [#:create-session? #t] [#:resource-limits '()]
  [#:environment-variables (default-environment-variables)]
fork+exec-command command [#:user #f] [#:group #f] [procedure]
  [#:supplementary-groups '()] [#:log-file #f] [#:log-encoding "UTF-8"]
  [#:directory (default-service-directory)] [#:file-creation-mask #f]
  [#:create-session? #t] [#:resource-limits '()] [#:environment-variables
  (default-environment-variables)]

```

Run *command* as the current process from *directory*, with *file-creation-mask* if it's true, with *rlimits*, and with *environment-variables* (a list of strings like "PATH=/bin".) File descriptors 1 and 2 are kept as is or redirected to either *log-port* or *log-file* if it's true, whereas file descriptor 0 (standard input) points to *input-port* or */dev/null*; all other file descriptors are closed prior to yielding control to *command*. When *create-session?* is true, call *setsid* first (see Section “Processes” in *GNU Guile Reference Manual*).

By default, *command* is run as the current user. If the *user* keyword argument is present and not false, change to *user* immediately before invoking *command*. *user* may be a string, indicating a user name, or a number, indicating a user ID. Likewise, *command* will be run under the current group, unless the *group* keyword argument is present and not false, and *supplementary-groups* is not '() .

fork+exec-command does the same as *exec-command*, but in a separate process whose PID it returns.

default-environment-variables [Scheme Variable]

This parameter (see Section “Parameters” in *GNU Guile Reference Manual*) specifies the default list of environment variables to be defined when the procedures above create a new process.

It must be a list of strings where each string has the format *name=value*. It defaults to what *environ* returns when the program starts (see Section “Runtime Environment” in *GNU Guile Reference Manual*).

default-pid-file-timeout [Scheme Variable]

This parameter (see Section “Parameters” in *GNU Guile Reference Manual*) specifies the default PID file timeout in seconds, when *#:pid-file* is used (see above). It defaults to 5 seconds.

One may also define services meant to be started *on demand*. In that case, shepherd listens for incoming connections on behalf of the program that handles them; when it accepts an incoming connection, it starts the program to handle them. The main benefit is that such services do not consume resources until they are actually used, and they do not slow down startup.

These services are implemented following the protocol of the venerable *inetd* “super server” (see Section “*inetd* invocation” in *GNU Inetutils*). Many network daemons can be invoked in “*inetd* mode”; this is the case, for instance, of *sshd*, the secure shell server of the OpenSSH project. The Shepherd lets you define *inetd*-style services, specifically those

in `nowait` mode where the daemon is passed the newly-accepted socket connection while `shepherd` is in charge of listening.

Listening endpoints for such services are described as records built using the `endpoint` procedure:

```
endpoint address [#:name "unknown"] [#:style SOCK_STREAM] [procedure]
  [#:backlog 128] [#:socket-owner (getuid)] [#:socket-group (getgid)]
  [#:socket-directory-permissions #o755]
```

Return a new endpoint called *name* of *address*, an address as return by `make-socket-address`, with the given *style* and *backlog*.

When *address* is of type `AF_INET6`, the endpoint is *IPv6-only*. Thus, if you want a service available both on IPv4 and IPv6, you need two endpoints. For example, below is a list of endpoints to listen on port 4444 on all the network interfaces, both in IPv4 and IPv6 (“0.0.0.0” for IPv4 and “::0” for IPv6):

```
(list (endpoint (make-socket-address AF_INET INADDR_ANY 4444))
      (endpoint (make-socket-address AF_INET6 IN6ADDR_ANY 4444)))
```

This is the list you would pass to `make-inetd-creator` or `make-systemd-creator`—see below.

When *address* is of type `AF_UNIX`, *socket-owner* and *socket-group* are strings or integers that specify its ownership and that of its parent directory; *socket-directory-permissions* specifies the permissions for its parent directory.

The `inetd` service constructor takes a command and a list of such endpoints:

```
make-inetd-creator command endpoints [procedure]
  [#:service-name-stem _] [#:requirements '()] [#:max-connections
  (default-inetd-max-connections)] [#:user #f] [#:group #f]
  [#:supplementary-groups '()] [#:directory (default-service-directory)]
  [#:file-creation-mask #f] [#:create-session? #t] [#:resource-limits '()]
  [#:environment-variables (default-environment-variables)]
```

Return a procedure that opens sockets listening to *endpoints*, a list of objects as returned by `endpoint`, and accepting connections in the background.

Upon a client connection, a transient service running *command* is spawned. Only up to *max-connections* simultaneous connections are accepted; when that threshold is reached, new connections are immediately closed.

The remaining arguments are as for `make-forkexec-creator`.

```
make-inetd-destructor [procedure]
  Return a procedure that terminates an inetd service.
```

The last type is *systemd-style services*. Like `inetd-style` services, those are started on demand when an incoming connection arrives, but using the protocol devised by the `systemd` service manager and referred to as *socket activation* (<https://www.freedesktop.org/software/systemd/man/daemon.html#Socket-Based%20Activation>). The main difference with `inetd-style` services is that `shepherd` hands over the listening socket(s) to the daemon; the daemon is then responsible for accepting incoming connections. A handful of

environment variables are set in the daemon’s execution environment (see below), which usually checks them using the `libsystemd` or `libelogind` client library helper functions (https://www.freedesktop.org/software/systemd/man/sd_listen_fds.html). The constructor and destructor for `systemd`-style daemons are described below.

```
make-systemd-creator command endpoints [#:user #f] [procedure]
  [#:group #f] [#:supplementary-groups '()] [#:directory
  (default-service-directory)] [#:file-creation-mask #f] [#:create-session?
  #t] [#:resource-limits '()] [#:environment-variables
  (default-environment-variables)]
```

Return a procedure that starts *command*, a program and list of argument, as a `systemd`-style service listening on *endpoints*, a list of `<endpoint>` objects.

command is started on demand on the first connection attempt on one of *endpoints*. It is passed the listening sockets for *endpoints* in file descriptors 3 and above; as such, it is equivalent to an `Accept=no` `systemd` socket unit (<https://www.freedesktop.org/software/systemd/man/systemd.socket.html>). The following environment variables are set in its environment:

`LISTEN_PID`

It is set to the PID of the newly spawned process.

`LISTEN_FDS`

It contains the number of sockets available starting from file descriptor 3—i.e., the length of *endpoints*.

`LISTEN_FDNames`

The colon-separated list of endpoint names.

This must be paired with `make-systemd-destructor`.

```
make-systemd-destructor [procedure]
  Return a procedure that terminates a systemd-style service as created by
  make-systemd-creator.
```

4.5 Service Examples

The configuration file of the `shepherd` command (see Section 3.1 [Invoking shepherd], page 5) defines, registers, and possibly starts *services*. Each service specifies other services it depends on and how it is started and stopped. The configuration file contains Scheme code that uses the programming interface of the (`shepherd service`) module (see Chapter 4 [Services], page 8).

Let’s assume you want to define and register a service to start `mcron`, the daemon that periodically executes jobs in the background (see Section “Introduction” in *GNU mcron Manual*). That service is started by invoking the `mcron` command, after which `shepherd` should monitor the running process, possibly re-spawning it if it stops unexpectedly. Here’s the configuration file for this one service:

```
(define mcron
  (make <service>
    #:provides '(mcron)
```

```
;; Run /usr/bin/mcron without any command-line arguments.
#:start (make-forkexec-constructor '("/usr/bin/mcron"))
#:stop (make-kill-destructor)
#:respawn? #t))
```

```
(register-services mcron)
```

You can write the snippet above in the default configuration file—`~/config/shepherd/init.scm` if you run `shepherd` as an unprivileged user. When you launch it, `shepherd` will evaluate that configuration; thus it will define and register the `mcron` service, but it will *not* start it. To start the service, run:

```
herd start mcron
```

Alternatively, if you want `mcron` to be started automatically when `shepherd` starts, you can add this snippet at the end of the configuration file:

```
(start-in-the-background '(mcron))
```

Now let's take another example: `sshd`, the secure shell daemon of the OpenSSH project (<https://www.openssh.com>). We will pass `sshd` the `-D` option so that it does not “detach”, making it easy for `shepherd` to monitor its process; we also tell `shepherd` to check its *PID file* to determine once it has started and is ready to accept connections:

```
(define sshd
  (make <service>
    #:provides '(sshd ssh-daemon) ;for convenience, give it two names
    #:start (make-forkexec-constructor
             '("/usr/sbin/sshd" "-D")
             #:pid-file "/etc/ssh/sshd.pid")
    #:stop (make-kill-destructor)
    #:respawn? #t))
```

```
(register-services sshd)
```

```
(start-in-the-background '(sshd))
```

Alternatively, we can start `sshd` in *inetd mode*: in that case, `shepherd` listens for connection and spawns `sshd` only upon incoming connections. The `inetd` mode is enabled by passing the `-i` command-line option:

```
(define sshd
  (make <service>
    #:provides '(sshd ssh-daemon)
    #:start (make-inetd-constructor
             '("/usr/sbin/sshd" "-D" "-i")
             (list (endpoint
                   (make-socket-address AF_INET INADDR_ANY 22))
                  (endpoint
                   (make-socket-address AF_INET6 IN6ADDR_ANY 22))))
             #:max-connections 10)
    #:stop (make-inetd-destructor)
    #:respawn? #t))
```

```
(register-services sshd)
(start-in-the-background '(sshd))
```

The `make-socket-address` procedure calls above return the *listening addresses* (see Section “Network Socket Address” in *GNU Guile Reference Manual*). In this case, it specifies that `shepherd` will accept connections coming from any network interface (“0.0.0.0” in IPv4 notation and “::0” for IPv6) on port 22. The `endpoint` calls wrap these addresses in endpoint records (see [endpoints], page 14). When a client connects, `shepherd` accepts it and spawns `sshd -D -i` as a new *transient service*, passing it the client connection. The `#:max-connections` parameter instructs `shepherd` to accept at most 10 simultaneous client connections.

In these examples, we haven’t discussed dependencies among services—the `#:requires` keyword of `<service>`—nor did we discuss `systemd`-style services. These are extensions of what we’ve seen so far. See Chapter 4 [Services], page 8, for details.

If you use Guix System, you will see that it contains a wealth of Shepherd service definitions. The nice thing is that those give you a *complete view* of what goes into the service—not just how the service is started and stopped, but also what software package is used and what configuration file is provided. See Section “Shepherd Services” in *GNU Guix Reference Manual*, for more info.

4.6 Managing User Services

The Shepherd can be used to manage services for an unprivileged user. First, you may want to ensure it is up and running every time you log in. One way to accomplish that is by adding the following lines to `~/.bash_profile` (see Section “Bash Startup Files” in *The GNU Bash Reference Manual*):

```
if [[ ! -S ${XDG_RUNTIME_DIR}-${HOME}/.cache}/shepherd/socket ]]; then
  shepherd
fi
```

Then, we suggest the following top-level `$XDG_CONFIG_HOME/shepherd/init.scm` file, which will automatically load individual service definitions from `~/.config/shepherd/init.d`:

```
(use-modules (shepherd service)
             ((ice-9 ftw) #:select (scandir)))

;; Load all the files in the directory 'init.d' with a suffix '.scm'.
(for-each
 (lambda (file)
  (load (string-append "init.d/" file)))
 (scandir (string-append (dirname (current-filename)) "/init.d")
  (lambda (file)
   (string-suffix? ".scm" file))))

;; Send shepherd into the background
(action 'shepherd 'daemonize)
```

Then, individual user services can be put in `$XDG_CONFIG_HOME/shepherd/init.d/`, e.g., for `ssh-agent`.

```
;;; Commentary:
```

```

;;;
;;; Add to your ~/.bash_profile:
;;;
;;; SSH_AUTH_SOCK=${XDG_RUNTIME_DIR-$HOME/.cache}/ssh-agent/socket
;;; export SSH_AUTH_SOCK
;;;
;;; Code:

(use-modules (shepherd support))

(define ssh-agent
  (make <service>
    #:provides '(ssh-agent)
    #:docstring "Run 'ssh-agent'"
    #:start (lambda ()
              (let ((socket-dir (string-append %user-runtime-dir "/ssh-agent")))
                (unless (file-exists? socket-dir)
                  (mkdir-p socket-dir)
                  (chmod socket-dir #o700))
                (fork+exec-command
                 ("ssh-agent" "-D" "-a" ,(string-append socket-dir "/socket"))
                 #:log-file (string-append %user-log-dir "/ssh-agent.log"))))
            #:stop (make-kill-destructor)
            #:respawn? #t))

(register-services ssh-agent)
(start ssh-agent)

```

4.7 The root and unknown services

The service `root` is special, because it is used to control the Shepherd itself. It has an alias `shepherd`. It provides the following actions (in addition to `enable`, `disable` and `restart` which do not make sense here).

status Displays which services are started and which ones are not.

detailed-status

Displays detailed information about every registered service.

load file Evaluate the Scheme code in *file* in a fresh module that uses the `(oop goops)` and `(shepherd services)` modules—as with the `--config` option of `shepherd` (see Section 3.1 [Invoking shepherd], page 5).

eval exp Likewise, evaluate Scheme expression *exp* in a fresh module with all the necessary bindings.

unload service-name

Attempt to remove the service identified by *service-name*. `shepherd` will first stop the service, if necessary, and then remove it from the list of registered

services. Any services depending upon *service-name* will be stopped as part of this process.

If *service-name* simply does not exist, output a warning and do nothing. If it exists, but is provided by several services, output a warning and do nothing. This latter case might occur for instance with the fictional service `web-server`, which might be provided by both `apache` and `nginx`. If *service-name* is the special string and `all`, attempt to remove all services except for the Shepherd itself.

reload *file-name*

Unload all known optional services using `unload`'s `all` option, then load *file-name* using `load` functionality. If *file-name* does not exist or `load` encounters an error, you may end up with no defined services. As these can be reloaded at a later stage this is not considered a problem. If the `unload` stage fails, `reload` will not attempt to load *file-name*.

daemonize

Fork and go into the background. This should be called before respawnable services are started, as otherwise we would not get the `SIGCHLD` signals when they terminate.

The `unknown` service must be defined by the user and if it exists, is used as a fallback whenever we try to invoke an unknown action of an existing service or use a service that does not exist. This is useful only in few cases, but enables you to do various sorts of unusual things.

5 Misc Facilities

This is a list of facilities which are available to code running inside of the Shepherd and is considered generally useful, but is not directly related to one of the other topic covered in this manual.

5.1 Errors

`assert expr` [macro]
 If *expr* yields `#f`, display an appropriate error message and throw an `assertion-failed` exception.

`caught-error key args` [procedure]
 Tell the Shepherd that a *key* error with *args* has occurred. This is the simplest way to cause caught error result in uniformly formatted warning messages. The current implementation is not very good, though.

`without-system-error expr . . .` [macro]
 Evaluates the *exprs*, not going further if a system error occurs, but also doing nothing about it.

5.2 Communication

The (`shepherd comm`) module provides primitives that allow clients such as `herd` to connect to `shepherd` and send it commands to control or change its behavior (see Section 4.1 [Slots of services], page 8).

Currently, clients may only send *commands*, represented by the `<shepherd-command>` type. Each command specifies a service it applies to, an action name, a list of strings to be used as arguments, and a working directory. Commands are instantiated with `shepherd-command`:

`shepherd-command action service [#:arguments '()]` [procedure]
`[#:directory (getcwd)]`
 Return a new command (a `<shepherd-command>`) object for *action* on *service*.

Commands may then be written to or read from a communication channel with the following procedures:

`write-command command port` [procedure]
 Write *command* to *port*.

`read-command port` [procedure]
 Receive a command from *port* and return it.

In practice, communication with `shepherd` takes place over a Unix-domain socket, as discussed earlier (see Section 3.1 [Invoking shepherd], page 5). Clients may open a connection with the procedure below.

`open-connection` [*file*] [procedure]

Open a connection to the daemon, using the Unix-domain socket at *file*, and return the socket.

When *file* is omitted, the default socket is used.

The daemon writes output to be logged or passed to the currently-connected client using `local-output`:

`local-output` *format-string* . *args* [procedure]

This procedure should be used for all output operations in the Shepherd. It outputs the *args* according to the *format-string*, then inserts a newline. It writes to whatever is the main output target of the Shepherd, which might be multiple at the same time in future versions.

Under the hood, `write-command` and `read-command` write/read commands as s-expressions (sexps). Each sexp is intelligible and specifies a protocol version. The idea is that users can write their own clients rather than having to invoke `herd`. For instance, when you type `herd status`, what is sent over the wire is the following sexp:

```
(shepherd-command
 (version 0)
 (action status) (service root)
 (arguments ()) (directory "/data/src/dmd"))
```

The reply is also an sexp, along these lines:

```
(reply (version 0)
 (result (((service ...) ...)))
 (error #f) (messages ()))
```

This reply indicates that the `status` action was successful, because `error` is `#f`, and gives a list of sexps denoting the status of services as its `result`. The `messages` field is a possibly-empty list of strings meant to be displayed as is to the user.

5.3 Others

`copy-hashq-table` *table* *new-size* [procedure]

Create a hash-table with size *new-size*, and insert all values from *table* into it, using `eq?` when inserting. This procedure is mainly used internally, but is a generally useful utility, so it can be used by everyone.

6 Internals

This chapter contains information about the design and the implementation details of the Shepherd for people who want to hack it.

The GNU Shepherd is developed by a group of people in connection with Guix System (<https://www.gnu.org/software/guix/>), GNU's advanced distribution, but it can be used on other distros as well. You're very much welcome to join us! You can report bugs to bug-guix@gnu.org and send patches or suggestions to guix-devel@gnu.org.

6.1 Coding standards

About formatting: Use common sense and GNU Emacs (which actually is the same, of course), and you almost can't get the formatting wrong. Formatting should be as in Guile and Guix, basically. See Section "Coding Style" in *GNU Guix Reference Manual*, for more info.

6.2 Design decisions

Note: This section was written by Wolfgang Jähring back in 2003 and documents the original design of what was then known as GNU dmd. The main ideas remain valid but some implementation details and goals have changed.

The general idea of a service manager that uses dependencies, similar to those of a Makefile, came from the developers of the GNU Hurd, but as few people are satisfied with System V Init, many other people had the same idea independently. Nevertheless, the Shepherd was written with the goal of becoming a replacement for System V Init on GNU/Hurd, which was one of the reasons for choosing the extension language of the GNU project, Guile, for implementation (another reason being that it makes it just so much easier).

The runlevel concept (i.e. thinking in *groups* of services) is sometimes useful, but often one also wants to operate on single services. System V Init makes this hard: While you can start and stop a service, `init` will not know about it, and use the runlevel configuration as its source of information, opening the door for inconsistencies (which fortunately are not a practical problem usually). In the Shepherd, this was avoided by having a central entity that is responsible for starting and stopping the services, which therefore knows which services are actually started (if not completely improperly used, but that is a requirement which is impossible to avoid anyway). While runlevels are not implemented yet, it is clear that they will sit on top of the service concept, i.e. runlevels will merely be an optional extension that the service concept does not rely on. This also makes changes in the runlevel design easier when it may become necessary.

The consequence of having a daemon running that controls the services is that we need another program as user interface which communicates with the daemon. Fortunately, this makes the commands necessary for controlling services pretty short and intuitive, and gives the additional bonus of adding some more flexibility. For example, it is easily possible to grant password-protected control over certain services to unprivileged users, if desired.

An essential aspect of the design of the Shepherd (which was already mentioned above) is that it should always know exactly what is happening, i.e. which services are started and

stopped. The alternative would have been to not use a daemon, but to save the state on the file system, again opening the door for inconsistencies of all sorts. Also, we would have to use a separate program for respawning a service (which just starts the services, waits until it terminates and then starts it again). Killing the program that does the respawning (but not the service that is supposed to be respawned) would cause horrible confusion. My understanding of “The Right Thing” is that this conceptionally limited strategy is exactly what we do not want.

The way dependencies work in the Shepherd took a while to mature, as it was not easy to figure out what is appropriate. I decided to not make it too sophisticated by trying to guess what the user might want just to theoretically fulfill the request we are processing. If something goes wrong, it is usually better to tell the user about the problem and let her fix it, taking care to make finding solutions or workarounds for problems (like a misconfigured service) easy. This way, the user is in control of what happens and we can keep the implementation simple. To make a long story short, *we don't try to be too clever*, which is usually a good idea in developing software.

If you wonder why I was giving a “misconfigured service” as an example above, consider the following situation, which actually is a wonderful example for what was said in the previous paragraph: Service X depends on symbol S, which is provided by both A and B. A depends on AA, B depends on BB. AA and BB conflict with each other. The configuration of A contains an error, which will prevent it from starting; no service is running, but we want to start X now. In resolving its dependencies, we first try to start A, which will cause AA to be started. After this is done, the attempt of starting A fails, so we go on to B, but its dependency BB will fail to start because it conflicts with the running service AA. So we fail to provide S, thus X cannot be started. There are several possibilities to deal with this:

- When starting A fails, terminate those services which have been started in order to fulfill its dependencies (directly and indirectly). In case AA was running already, we would not want to terminate it. Well, maybe we would, to avoid the conflict with BB. But even if we would find out somehow that we need to terminate AA to eventually start X, is the user aware of this and wants this to happen (assuming AA was running already)? Probably not, she very likely has assumed that starting A succeeds and thus terminating AA is not necessary. Remember, unrelated (running) services might depend in AA. Even if we ignore this issue, this strategy is not only complicated, but also far from being perfect: Let's assume starting A succeeds, but X also depends on a service Z, which requires BB. In that case, we would need to detect in the first place that we should not even try to start A, but directly satisfy X's dependency on S with B.
- We could do it like stated above, but stop AA only if we know we won't need it anymore (for resolving further dependencies), and start it only when it does not conflict with anything that needs to get started. But should we stop it if it conflicts with something that *might* get started? (We do not always know for sure what we will start, as starting a service might fail and we want to fall back to a service that also provides the particular required symbol in that case.) I think that either decision will be bad in one case or another, even if this solution is already horribly complicated.
- When we are at it, we could just calculate a desired end-position, and try to get there by starting (and stopping!) services, recalculating what needs to be done whenever starting a service fails, also marking that particular service as unstartable, except if it

fails to start because a dependency could not be resolved (or maybe even then?). This is even more complicated. Instead of implementing this and thereby producing code that (a) nobody understands, (b) certainly has a lot of bugs, (c) will be unmaintainable and (d) causes users to panic because they won't understand what will happen, I decided to do the following instead:

- Just report the error, and let the user fix it (in this case, fix the configuration of A) or work around it (in this case, disable A so that we won't start AA but directly go on to starting B).

I hope you can agree that the latter solution after all is the best one, because we can be sure to not do something that the user does not want us to do. Software should not run amok. This explanation was very long, but I think it was necessary to justify why the Shepherd uses a very primitive algorithm to resolve dependencies, despite the fact that it could theoretically be a bit more clever in certain situations.

One might argue that it is possible to ask the user if the planned actions are ok with her, and if the plan changes ask again, but especially given that services are supposed to usually work, I see few reasons to make the source code of the Shepherd more complicated than necessary. If you volunteer to write *and* maintain a more clever strategy (and volunteer to explain it to everyone who wants to understand it), you are welcome to do so, of course. . .

6.3 Service Internals

Appendix A GNU Free Documentation License

Version 1.3, 3 November 2008

Copyright © 2000, 2001, 2002, 2007, 2008 Free Software Foundation, Inc.
<http://fsf.org/>

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document *free* in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or non-commercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of “copyleft”, which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The “Document”, below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as “you”. You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A “Modified Version” of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A “Secondary Section” is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document’s overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The “Invariant Sections” are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released

under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The “Cover Texts” are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A “Transparent” copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not “Transparent” is called “Opaque”.

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The “Title Page” means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, “Title Page” means the text near the most prominent appearance of the work’s title, preceding the beginning of the body of the text.

The “publisher” means any person or entity that distributes copies of the Document to the public.

A section “Entitled XYZ” means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as “Acknowledgements”, “Dedications”, “Endorsements”, or “History”.) To “Preserve the Title” of such a section when you modify the Document means that it remains a section “Entitled XYZ” according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any,

- be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
 - C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
 - D. Preserve all the copyright notices of the Document.
 - E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
 - F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
 - G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
 - H. Include an unaltered copy of this License.
 - I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
 - J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
 - K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
 - L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
 - M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
 - N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
 - O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their

titles to the list of Invariant Sections in the Modified Version’s license notice. These titles must be distinct from any other section titles.

You may add a section Entitled “Endorsements”, provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled “History” in the various original documents, forming one section Entitled “History”; likewise combine any sections Entitled “Acknowledgements”, and any sections Entitled “Dedications”. You must delete all sections Entitled “Endorsements.”

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an “aggregate” if the copyright resulting from the compilation is not used to limit the legal rights of the compilation’s users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document’s Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled “Acknowledgements”, “Dedications”, or “History”, the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, or distribute it is void, and will automatically terminate your rights under this License.

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, receipt of a copy of some or all of the same material does not give you any rights to use it.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License “or any later version” applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation. If the Document specifies that a proxy can decide which future versions of this License can be used, that proxy’s public statement of acceptance of a version permanently authorizes you to choose that version for the Document.

11. RELICENSING

“Massive Multiauthor Collaboration Site” (or “MMC Site”) means any World Wide Web server that publishes copyrightable works and also provides prominent facilities for anybody to edit those works. A public wiki that anybody can edit is an example of such a server. A “Massive Multiauthor Collaboration” (or “MMC”) contained in the site means any set of copyrightable works thus published on the MMC site.

“CC-BY-SA” means the Creative Commons Attribution-Share Alike 3.0 license published by Creative Commons Corporation, a not-for-profit corporation with a principal place of business in San Francisco, California, as well as future copyleft versions of that license published by that same organization.

“Incorporate” means to publish or republish a Document, in whole or in part, as part of another Document.

An MMC is “eligible for relicensing” if it is licensed under this License, and if all works that were first published under this License somewhere other than this MMC, and subsequently incorporated in whole or in part into the MMC, (1) had no cover texts or invariant sections, and (2) were thus incorporated prior to November 1, 2008.

The operator of an MMC Site may republish an MMC contained in the site under CC-BY-SA on the same site at any time before August 1, 2009, provided the MMC is eligible for relicensing.

ADDENDUM: How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

```
Copyright (C) year your name.  
Permission is granted to copy, distribute and/or modify this document  
under the terms of the GNU Free Documentation License, Version 1.3  
or any later version published by the Free Software Foundation;  
with no Invariant Sections, no Front-Cover Texts, and no Back-Cover  
Texts. A copy of the license is included in the section entitled ‘GNU  
Free Documentation License’.
```

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the “with...Texts.” line with this:

```
with the Invariant Sections being list their titles, with  
the Front-Cover Texts being list, and with the Back-Cover Texts  
being list.
```

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.

Concept Index

- <
- <service>, slots of 8
- ## A
- Actions of services 9
- assertions 20
- ## C
- canonical names of services 8
- Configuration file 2
- constructors, generation of 11
- ## D
- daemon 5
- daemon controller 5
- deco, daemon controller 5
- destructors, generation of 11
- dmd 1
- ## E
- endpoints, for inetd services 14
- endpoints, for systemd services 14
- ## F
- fallback service 19
- fallback services 3
- ## G
- generating constructors 11
- generating destructors 11
- GOOPS 1
- Guile 1
- ## H
- hashes 21
- herd 5, 6, 7
- herding, of daemons 5
- Hook for individual services 8
- ## I
- inetd mode, example 16
- inetd-style services 13
- insecure 5
- invoking **shepherd** 5
- ## L
- log file 5
- logging 5
- ## O
- on-demand, starting services 13, 14
- one-shot services 8
- output 21
- ## P
- prefix 2
- protocol, between **shepherd** and its clients 21
- ## R
- relative file names 5
- Respawning services 8
- root service 18
- ## S
- Scheme 1
- security 5
- service 8
- Service actions 9
- Service constructor 9
- Service destructor 9
- service manager 1
- Service status 2
- shepherd 5
- shepherd** Invocation 5
- slots of <service> 8
- socket activation, starting services 14
- socket special file 6
- special services 18
- Starting a service 2, 9
- starting services, via socket activation 14
- Status (of services) 2
- Stopping a service 2, 9
- syslog 5
- system errors 20
- systemd-style services 14
- ## T
- Termination of a service's process 9
- transient services 8
- ## U
- unknown service 19

V

virtual services 3

Procedure and Macro Index

A

action 10, 11
 assert 20

C

canonical-name 10
 caught-error 20
 conflicts-with 10
 copy-hashq-table 21

D

default-display-status 10

E

endpoint 14
 exec-command 13

F

find-running 11
 for-each-service 11
 fork+exec-command 13

L

local-output 21
 lookup-services 11

M

make-actions 11
 make-forkexec-creator 12
 make-inetd-creator 14
 make-inetd-destructor 14
 make-kill-destructor 12
 make-system-creator 12
 make-system-destructor 12
 make-systemd-creator 15
 make-systemd-destructor 15

O

one-shot? 10
 open-connection 21

P

provided-by 10

R

read-command 20
 register-services 11
 required-by 10
 respawn? 10
 running? 10

S

shepherd-command 20
 start 10, 11
 start-in-the-background 11
 stop 10, 11

W

without-system-error 20
 write-command 20

Variable Index

A

actions (slot of <service>) 9

D

default-environment-variables 13

default-pid-file-timeout 13

E

enabled? (slot of <service>) 9

H

handle-termination (slot of <service>) 9

L

last-respawns (slot of <service>) 9

O

one-shot? (slot of <service>) 8

P

provides (slot of <service>) 8

R

replacement (slot of <service>) 9

requires (slot of <service>) 8

respawn? (slot of <service>) 8

running (slot of <service>) 8

S

start (slot of <service>) 9

stop (slot of <service>) 9

stop-delay? (slot of <service>) 9

T

transient? (slot of <service>) 8

W

waiting-for-termination? (slot
of <service>) 9

X

XDG_RUNTIME_DIR 6

Type Index

<service>.....	8
<shepherd-command>.....	20