ASDF: Another System Definition Facility

This manual describes ASDF, a system definition facility for Common Lisp programs and libraries.

You can find the latest version of this manual at http://common-lisp.net/project/asdf/asdf.html.

ASDF Copyright © 2001-2010 Daniel Barlow and contributors.

This manual Copyright © 2001-2010 Daniel Barlow and contributors.

This manual revised © 2009-2010 Robert P. Goldman and Francois-Rene Rideau.

Table of Contents

7   Controlling where ASDF searches for systems

8   Controlling where ASDF saves compiled files

9   Error handling

10   Miscellaneous additional functionality

11   Getting the latest version

1 Introduction

ASDF is Another System Definition Facility: a tool for specifying how systems of Common Lisp software are comprised of components (sub-systems and files), and how to operate on these components in the right order so that they can be compiled, loaded, tested, etc.

ASDF presents three faces: one for users of Common Lisp software who want to reuse other people's code, one for writers of Common Lisp software who want to specify how to build their systems, one for implementers of Common Lisp extensions who want to extend the build system. See Chapter 4 [Loading a system], page 7, to learn how to use ASDF to load a system. See Chapter 5 [Defining systems with defsystem], page 8, to learn how to define a system of your own. See Chapter 6 [The object model of ASDF], page 14, for a description of the ASDF internals and how to extend ASDF.

Nota Bene: We have released ASDF 2.000 on May 31st 2010. It hopefully will have been it included in all CL maintained implementations shortly afterwards. See Chapter 12 ["What has changed between ASDF 1 and ASDF 2?"], page 36.

2 Loading ASDF

2.1 Loading a pre-installed ASDF

Many Lisp implementations include a copy of ASDF. You can usually load this copy using Common Lisp's `require` function:

```
(require :asdf)
```

Consult your Lisp implementation's documentation for details.

Hopefully, ASDF 2 will soon be bundled with every Common Lisp implementation, and you can load it that way.

2.2 Checking whether ASDF is loaded

To check whether ASDF is properly loaded in your current Lisp image, you can run this form:

```
(asdf:asdf-version)
```

If it returns a string, that is the version of ASDF that is currently installed.

If it raises an error, then either ASDF is not loaded, or you are using an old version of ASDF.

You can check whether an old version is loaded by checking if the ASDF package is present. The form below will allow you to programmatically determine whether a recent version is loaded, an old version is loaded, or none at all:

```
(or #+asdf2 (asdf:asdf-version) #+asdf :old)
```

If it returns a version number, that's the version of ASDF installed. If it returns the keyword `:OLD`, then you're using an old version of ASDF (from before 1.635). If it returns `NIL` then ASDF is not installed.

If you are running a version older than 2.000, we recommend that you load a newer ASDF using the method below.

2.3 Upgrading ASDF

If your implementation does provide ASDF 2 or later, and you want to upgrade to a more recent version, just install ASDF like any other package (see see Chapter 2 [Loading an otherwise installed ASDF], page 2 below), configure ASDF as usual (see see Chapter 3 [Configuring ASDF], page 4 below), and upgrade with:

```
(require :asdf)
(asdf:load-system :asdf)
```

If on the other hand, your implementation only provides an old ASDF, you will require a special configuration step and an old-style loading:

```
(require :asdf)
(push #p"/path/to/new/asdf/" asdf:*central-registry*)
(asdf:oos 'asdf:load-op :asdf)
```

Don't forget the trailing `/` at the end of your pathname.

Also, note that older versions of ASDF won't redirect their output, or at least won't do it according to your usual ASDF 2 configuration. You therefore need write access on the directory where you install the new ASDF, and make sure you're not using it for multiple mutually incompatible implementations. At worst, you may have to have multiple copies of the new ASDF, e.g. one per implementation installation, to avoid clashes.

Finally, note that there are some limitations to upgrading ASDF:

- Any ASDF extension is invalidated, and will need to be reloaded.
- It is safer if you upgrade ASDF and its extensions as a special step at the very beginning of whatever script you are running, before you start using ASDF to load anything else.

2.4 Loading an otherwise installed ASDF

If your implementation doesn't include ASDF, if for some reason the upgrade somehow fails, does not or cannot apply to your case, you will have to install the file 'asdf.lisp' somewhere and load it with:

```
(load "/path/to/your/installed/asdf.lisp")
```

The single file 'asdf.lisp' is all you normally need to use ASDF.

You can extract this file from latest release tarball on the ASDF website. If you are daring and willing to report bugs, you can get the latest and greatest version of ASDF from its git repository. See Chapter 11 [Getting the latest version], page 35.

For maximum convenience you might want to have ASDF loaded whenever you start your Lisp implementation, for example by loading it from the startup script or dumping a custom core — check your Lisp implementation's manual for details.

3 Configuring ASDF

3.1 Configuring ASDF to find your systems

So it may compile and load your systems, ASDF must be configured to find the '`.asd`' files that contain system definitions.

Since ASDF 2, the preferred way to configure where ASDF finds your systems is the `source-registry` facility, fully described in its own chapter of this manual. See Chapter 7 [Controlling where ASDF searches for systems], page 21.

The default location for a user to install Common Lisp software is under '`~/.local/share/common-lisp/source/`'. If you install software there, you don't need further configuration. If you're installing software yourself at a location that isn't standard, you have to tell ASDF where you installed it. See below. If you're using some tool to install software, the authors of that tool should already have configured ASDF.

The simplest way to add a path to your search path, say '`/home/luser/.asd-link-farm/`' is to create the directory '`~/.config/common-lisp/source-registry.conf.d/`' and there create a file with any name of your choice but the type '`conf`', for instance '`42-asd-link-farm.conf`' containing the line:

```
(:directory "/home/luser/.asd-link-farm/")
```

If you want all the subdirectories under '`/home/luser/lisp/`' to be recursively scanned for '`.asd`' files, instead use:

```
(:tree "/home/luser/lisp/")
```

Note that your Operating System distribution or your system administrator may already have configured system-managed libraries for you.

The required '`.conf`' extension allows you to have disabled files or editor backups (ending in '`~`'), and works portably (for instance, it is a pain to allow both empty and non-empty extension on CLISP). Excluded are files the name of which start with a '`.`' character. It is customary to start the filename with two digits that specify the order in which the directories will be scanned.

ASDF will automatically read your configuration the first time you try to find a system. You can reset the source-registry configuration with:

```
(asdf:clear-source-registry)
```

And you probably should do so before you dump your Lisp image, if the configuration may change between the machine where you save it at the time you save it and the machine you resume it at the time you resume it.

3.2 Configuring ASDF to find your systems – old style

The old way to configure ASDF to find your systems is by `push`ing directory pathnames onto the variable `asdf:*central-registry*`.

You must configure this variable between the time you load ASDF and the time you first try to use it. Loading and configuring ASDF presumably happen as part of some initialization script that builds or starts your Common Lisp software system. (For instance, some SBCL users used to put it in their '`~/.sbclrc`'.)

The `asdf:*central-registry*` is empty by default in ASDF 2, but is still supported for compatibility with ASDF 1. When used, it takes precedence over the above source-registry[1].

For instance, if you wanted ASDF to find the '.asd' file '/home/me/src/foo/foo.asd' your initialization script could after it loads ASDF with `(require :asdf)` configure it with:

```
(push "/home/me/src/foo/" asdf:*central-registry*)
```

Note the trailing slash: when searching for a system, ASDF will evaluate each entry of the central registry and coerce the result to a pathname[2] at which point the presence of the trailing directory name separator is necessary to tell Lisp that you're discussing a directory rather than a file.

Typically, however, there are a lot of '.asd' files, and a common idiom was to have to put a bunch of *symbolic links* to '.asd' files in a common directory and push *that* directory (the "link farm") to the `asdf:*central-registry*` instead of pushing each of the many involved directories to the `asdf:*central-registry*`. ASDF knows how to follow such *symlinks* to the actual file location when resolving the paths of system components (on Windows, you can use Windows shortcuts instead of POSIX symlinks).

For example, if `#p"/home/me/cl/systems/"` (note the trailing slash) is a member of `*central-registry*`, you could set up the system *foo* for loading with asdf with the following commands at the shell:

```
$ cd /home/me/cl/systems/
$ ln -s ~/src/foo/foo.asd .
```

This old style for configuring ASDF is not recommended for new users, but it is supported for old users, and for users who want to programmatically control what directories are added to the ASDF search path.

3.3 Configuring where ASDF stores object files

ASDF lets you configure where object files will be stored. Sensible defaults are provided and you shouldn't normally have to worry about it.

This allows the same source code repository may be shared between several versions of several Common Lisp implementations, between several users using different compilation options and without write privileges on shared source directories, etc. This also allows to keep source directories uncluttered by plenty of object files.

Starting with ASDF 2, the `asdf-output-translations` facility was added to ASDF itself, that controls where object files will be stored. This facility is fully described in a chapter of this manual, Chapter 8 [Controlling where ASDF saves compiled files], page 27.

---

[1] It is possible to further customize the system definition file search. That's considered advanced use, and covered later: search forward for `*system-definition-search-functions*`. See Chapter 5 [Defining systems with defsystem], page 8.

[2] ASDF will indeed call `EVAL` on each entry. It will also skip entries that evaluate to `NIL`.

Strings and pathname objects are self-evaluating, in which case the `EVAL` step does nothing; but you may push arbitrary SEXP onto the central registry, that will be evaluated to compute e.g. things that depend on the value of shell variables or the identity of the user.

The variable `asdf:*central-registry*` is thus a list of "system directory designators". A *system directory designator* is a form which will be evaluated whenever a system is to be found, and must evaluate to a directory to look in. By "directory" here, we mean "designator for a pathname with a supplied DIRECTORY component".

The simplest way to add a translation to your search path, say from '/foo/bar/baz/quux/' to '/where/i/want/my/fasls/' is to create the directory '~/.config/common-lisp/asdf-output-translations.conf.d/' and there create a file with any name of your choice and the type 'conf', for instance '42-bazquux.conf' containing the line:

```
("/foo/bar/baz/quux/" "/where/i/want/my/fasls/")
```

To disable output translations for source under a given directory, say '/toto/tata/' you can create a file '40-disable-toto.conf' with the line:

```
("/toto/tata/")
```

To wholly disable output translations for all directories, you can create a file '00-disable.conf' with the line:

```
(t t)
```

Note that your Operating System distribution or your system administrator may already have configured translations for you. In absence of any configuration, the default is to redirect everything under an implementation-dependent subdirectory of '~/.cache/common-lisp/'. See Chapter 7 [Controlling where ASDF searches for systems], page 21, for full details.

The required '.conf' extension allows you to have disabled files or editor backups (ending in '~'), and works portably (for instance, it is a pain to allow both empty and non-empty extension on CLISP). Excluded are files the name of which start with a '.' character. It is customary to start the filename with two digits that specify the order in which the directories will be scanned.

ASDF will automatically read your configuration the first time you try to find a system. You can reset the source-registry configuration with:

```
(asdf:clear-output-translations)
```

And you probably should do so before you dump your Lisp image, if the configuration may change between the machine where you save it at the time you save it and the machine you resume it at the time you resume it.

Finally note that before ASDF 2, other ASDF add-ons offered the same functionality, each in subtly different and incompatible ways: ASDF-Binary-Locations, cl-launch, common-lisp-controller. ASDF-Binary-Locations is now not needed anymore and should not be used. cl-launch 3.000 and common-lisp-controller 7.2 have been updated to just delegate this functionality to ASDF.

4 Using ASDF

4.1 Loading a system

The system *foo* is loaded (and compiled, if necessary) by evaluating the following Lisp form:

```
(asdf:load-system :foo)
```

On some implementations (namely ABCL, Clozure CL, CMUCL, ECL and SBCL), ASDF hooks into the `CL:REQUIRE` facility and you can just use:

```
(require :foo)
```

In older versions of ASDF, you needed to use (`asdf:oos 'asdf:load-op :foo`). If your ASDF is too old to provide `asdf:load-system` though we recommend that you upgrade to ASDF 2. See Chapter 2 [Loading an otherwise installed ASDF], page 2.

4.2 Other Operations

ASDF provides three commands for the most common system operations: `load-system`, `compile-system` or `test-system`.

Because ASDF is an extensible system for defining *operations* on *components*, it also provides a generic function `operate` (which is usually abbreviated by `oos`). You'll use `oos` whenever you want to do something beyond compiling, loading and testing.

Output from ASDF and ASDF extensions are supposed to be sent to the CL stream `*standard-output*`, and so rebinding that stream around calls to `asdf:operate` should redirect all output from ASDF operations.

Reminder: before ASDF can operate on a system, however, it must be able to find and load that system's definition. See Chapter 3 [Configuring ASDF to find your systems], page 4.

4.3 Summary

To use ASDF:

- Load ASDF itself into your Lisp image, either through (`require :asdf`) or else through (`load "/path/to/asdf.lisp"`).
- Make sure ASDF can find system definitions thanks to proper source-registry configuration.
- Load a system with (`load-system :my-system`) or use some other operation on some system of your choice.

4.4 Moving on

That's all you need to know to use ASDF to load systems written by others. The rest of this manual deals with writing system definitions for Common Lisp software you write yourself, including how to extend ASDF to define new operation and component types.

5 Defining systems with defsystem

This chapter describes how to use asdf to define systems and develop software.

5.1 The defsystem form

Systems can be constructed programmatically by instantiating components using `make-instance`. Most of the time, however, it is much more practical to use a static `defsystem` form. This section begins with an example of a system definition, then gives the full grammar of `defsystem`.

Let's look at a simple system. This is a complete file that would usually be saved as 'hello-lisp.asd':

```
(in-package :asdf)

(defsystem "hello-lisp"
  :description "hello-lisp: a sample Lisp system."
  :version "0.2"
  :author "Joe User <joe@example.com>"
  :licence "Public Domain"
  :components ((:file "packages")
               (:file "macros" :depends-on ("packages"))
               (:file "hello" :depends-on ("macros"))))
```

Some notes about this example:

- The file starts with an `in-package` form to use package `asdf`. You could instead start your definition by using a qualified name `asdf:defsystem`.

- If in addition to simply using `defsystem`, you are going to define functions, create ASDF extension, globally bind symbols, etc., it is recommended that to avoid namespace pollution between systems, you should create your own package for that purpose, for instance replacing the above (`in-package :asdf`) with:

  ```
  (defpackage :foo-system
    (:use :cl :asdf))

  (in-package :foo-system)
  ```

- The `defsystem` form defines a system named `hello-lisp` that contains three source files: 'packages', 'macros' and 'hello'.

- The file 'macros' depends on 'packages' (presumably because the package it's in is defined in 'packages'), and the file 'hello' depends on 'macros' (and hence, transitively on 'packages'). This means that ASDF will compile and load 'packages' and 'macros' before starting the compilation of file 'hello'.

- The files are located in the same directory as the file with the system definition. ASDF resolves symbolic links (or Windows shortcuts) before loading the system definition file and stores its location in the resulting system[1]. This is a good thing because the user can move the system sources without having to edit the system definition.

---

[1]   It is possible, though almost never necessary, to override this behaviour.

5.2 A more involved example

Let's illustrate some more involved uses of `defsystem` via a slightly convoluted example:

```
(defsystem "foo"
  :version "1.0"
  :components ((:module "foo" :components ((:file "bar") (:file"baz")
                                           (:file "quux"))
               :perform (compile-op :after (op c)
                           (do-something c))
               :explain (compile-op :after (op c)
                           (explain-something c)))
              (:file "blah")))
```

The method-form tokens need explaining: essentially, this part:

```
               :perform (compile-op :after (op c)
                           (do-something c))
               :explain (compile-op :after (op c)
                           (explain-something c))
```

has the effect of

```
(defmethod perform :after ((op compile-op) (c (eql ...)))
           (do-something c))
(defmethod explain :after ((op compile-op) (c (eql ...)))
           (explain-something c))
```

where `...` is the component in question; note that although this also supports `:before` methods, they may not do what you want them to — a `:before` method on perform `((op compile-op) (c (eql ...)))` will run after all the dependencies and sub-components have been processed, but before the component in question has been compiled.

5.3 The defsystem grammar

```
system-definition := ( defsystem system-designator option* )

option := :components component-list
        | :pathname pathname-specifier
        | :default-component-class
        | :perform method-form
        | :explain method-form
        | :output-files method-form
        | :operation-done-p method-form
        | :depends-on ( dependency-def* )
        | :serial [ t | nil ]
        | :in-order-to ( dependency+ )

component-list := ( component-def* )

component-def  := ( component-type simple-component-name option* )
```

```
component-type := :system | :module | :file | :static-file | other-component-type

other-component-type := symbol-by-name (see Section 5.3 [Component types],
page 9)

dependency-def := simple-component-name
               | ( :feature name )
               | ( :version simple-component-name version-specifier)

dependency := (dependent-op requirement+)
requirement := (required-op required-component+)
             | (feature feature-name)
dependent-op := operation-name
required-op := operation-name | feature

simple-component-name := string
                       | symbol

pathname-specifier := pathname | string | symbol

method-form := (operation-name qual lambda-list &rest body)
qual := method qualifier
```

5.3.1 Component names

Component names (`simple-component-name`) may be either strings or symbols.

5.3.2 Component types

Component type names, even if expressed as keywords, will be looked up by name in the current package and in the asdf package, if not found in the current package. So a component type `my-component-type`, in the current package `my-system-asd` can be specified as `:my-component-type`, or `my-component-type`.

5.3.3 Pathname specifiers

A pathname specifier (`pathname-specifier`) may be a pathname, a string or a symbol. When no pathname specifier is given for a component, which is the usual case, the component name itself is used.

If a string is given, which is the usual case, the string will be interpreted as a Unix-style pathname where `/` characters will be interpreted as directory separators. Usually, Unix-style relative pathnames are used (i.e. not starting with `/`, as opposed to absolute pathnames); they are relative to the path of the parent component. Finally, depending on the `component-type`, the pathname may be interpreted as either a file or a directory, and if it's a file, a file type may be added corresponding to the `component-type`, or else it will be extracted from the string itself (if applicable).

For instance, the `component-type` `:module` wants a directory pathname, and so a string `"foo/bar"` will be interpreted as the pathname '`#p"foo/bar/"`'. On the other hand, the `component-type` `:file` wants a file of type `lisp`, and so a string `"foo/bar"` will be inter-

preted as the pathname '`#p"foo/bar.lisp"`', and a string `"foo/bar.quux"` will be interpreted as the pathname '`#p"foo/bar.quux.lisp"`'. Finally, the `component-type :static-file` wants a file without specifying a type, and so a string `"foo/bar"` will be interpreted as the pathname '`#p"foo/bar"`', and a string `"foo/bar.quux"` will be interpreted as the pathname '`#p"foo/bar.quux"`'.

ASDF does not interpret the string `".."` to designate the parent directory. This string will be passed through to the underlying operating system for interpretation. We *believe* that this will work on all platforms where ASDF is deployed, but do not guarantee this behavior. A pathname object with a relative directory component of `:up` or `:back` is the only guaranteed way to specify a parent directory.

If a symbol is given, it will be translated into a string, and downcased in the process. The downcasing of symbols is unconventional, but was selected after some consideration. Observations suggest that the type of systems we want to support either have lowercase as customary case (Unix, Mac, windows) or silently convert lowercase to uppercase (lpns), so this makes more sense than attempting to use `:case :common` as argument to `make-pathname`, which is reported not to work on some implementations.

Pathname objects may be given to override the path for a component. Such objects are typically specified using reader macros such as `#p` or `#.(make-pathname ...)`. Note however, that `#p...` is a shorthand for `#.(parse-namestring ...)` and that the behavior of `parse-namestring` is completely non-portable, unless you are using Common Lisp `logical-pathnames` (see Section 5.3 [Warning about logical pathnames], page 9, below). Pathnames made with `#.(make-pathname ...)` can usually be done more easily with the string syntax above. The only case that you really need a pathname object is to override the component-type default file type for a given component. Therefore, pathname objects should only rarely be used. Unhappily, ASDF 1 didn't properly support parsing component names as strings specifying paths with directories, and the cumbersome `#.(make-pathname ...)` syntax had to be used.

Note that when specifying pathname objects, ASDF does not do any special interpretation of the pathname influenced by the component type, unlike the procedure for pathname-specifying strings. On the one hand, you have to be careful to provide a pathname that correctly fulfills whatever constraints are required from that component type (e.g. naming a directory or a file with appropriate type); on the other hand, you can circumvent the file type that would otherwise be forced upon you if you were specifying a string.

5.3.4 Warning about logical pathnames

We recommend that you not use logical pathnames in your asdf system definitions at this point, but logical pathnames *are* supported.

To use logical pathnames, you will have to provide a pathname object as a `:pathname` specifier to components that use it, using such syntax as `#p"LOGICAL-HOST:absolute;path;to;component.lisp"`.

You only have to specify such logical pathname for your system or some top-level component. Sub-components' relative pathnames, specified using the string syntax for names, will be properly merged with the pathnames of their parents. The specification of a logical pathname host however is *not* otherwise directly supported in the ASDF syntax for pathname specifiers as strings.

The `asdf-output-translation` layer will avoid trying to resolve and translate logical-pathnames. The advantage of this is that you can define yourself what translations you want to use with the logical pathname facility. The disadvantage is that if you do not define such translations, any system that uses logical pathnames will be have differently under asdf-output-translations than other systems you use.

If you wish to use logical pathnames you will have to configure the translations yourself before they may be used. ASDF currently provides no specific support for defining logical pathname translations.

5.3.5 Serial dependencies

If the `:serial t` option is specified for a module, ASDF will add dependencies for each each child component, on all the children textually preceding it. This is done as if by `:depends-on`.

```
:serial t
:components ((:file "a") (:file "b") (:file "c"))
```

is equivalent to

```
:components ((:file "a")
             (:file "b" :depends-on ("a"))
             (:file "c" :depends-on ("a" "b")))
```

5.3.6 Source location

The `:pathname` option is optional in all cases for systems defined via `defsystem`, and in the usual case the user is recommended not to supply it.

Instead, ASDF follows a hairy set of rules that are designed so that

1. `find-system` will load a system from disk and have its pathname default to the right place.

2. This pathname information will not be overwritten with `*default-pathname-defaults*` (which could be somewhere else altogether) if the user loads up the '`.asd`' file into his editor and interactively re-evaluates that form.

If a system is being loaded for the first time, its top-level pathname will be set to:

- The host/device/directory parts of `*load-truename*`, if it is bound.
- `*default-pathname-defaults*`, otherwise.

If a system is being redefined, the top-level pathname will be

- changed, if explicitly supplied or obtained from `*load-truename*` (so that an updated source location is reflected in the system definition)
- changed if it had previously been set from `*default-pathname-defaults*`
- left as before, if it had previously been set from `*load-truename*` and `*load-truename*` is currently unbound (so that a developer can evaluate a `defsystem` form from within an editor without clobbering its source location)

5.4 Other code in .asd files

Files containing `defsystem` forms are regular Lisp files that are executed by `load`. Consequently, you can put whatever Lisp code you like into these files (e.g., code that examines the compile-time environment and adds appropriate features to `*features*`). However, some conventions should be followed, so that users can control certain details of execution of the Lisp in '`.asd`' files:

- Any informative output (other than warnings and errors, which are the condition system's to dispose of) should be sent to the standard CL stream `*standard-output*`, so that users can easily control the disposition of output from ASDF operations.

6 The object model of ASDF

ASDF is designed in an object-oriented way from the ground up. Both a system's structure and the operations that can be performed on systems follow a protocol. ASDF is extensible to new operations and to new component types. This allows the addition of behaviours: for example, a new component could be added for Java JAR archives, and methods specialised on `compile-op` added for it that would accomplish the relevant actions.

This chapter deals with *components*, the building blocks of a system, and *operations*, the actions that can be performed on a system.

6.1 Operations

An *operation* object of the appropriate type is instantiated whenever the user wants to do something with a system like

- compile all its files
- load the files into a running lisp environment
- copy its source files somewhere else

Operations can be invoked directly, or examined to see what their effects would be without performing them. *FIXME: document how!* There are a bunch of methods specialised on operation and component type that actually do the grunt work.

The operation object contains whatever state is relevant for this purpose (perhaps a list of visited nodes, for example) but primarily is a nice thing to specialise operation methods on and easier than having them all be `EQL` methods.

Operations are invoked on systems via `operate`.

operate *operation system* &rest *initargs*   [Generic function]
oos *operation system* &rest *initargs*   [Generic function]
> `operate` invokes *operation* on *system*. `oos` is a synonym for `operate`.
>
> *operation* is a symbol that is passed, along with the supplied *initargs*, to `make-instance` to create the operation object. *system* is a system designator.
>
> The *initargs* are passed to the `make-instance` call when creating the operation object. Note that dependencies may cause the operation to invoke other operations on the system or its components: the new operations will be created with the same *initargs* as the original one.

6.1.1 Predefined operations of ASDF

All the operations described in this section are in the `asdf` package. They are invoked via the `operate` generic function.

    (asdf:operate 'asdf:operation-name :system-name {operation-options ...})

compile-op &key proclamations   [Operation]
> This operation compiles the specified component. If proclamations are supplied, they will be proclaimed. This is a good place to specify optimization settings.
>
> When creating a new component type, you should provide methods for `compile-op`.

When `compile-op` is invoked, component dependencies often cause some parts of the system to be loaded as well as compiled. Invoking `compile-op` does not necessarily load all the parts of the system, though; use `load-op` to load a system.

`load-op` *&key* `proclamations`   [Operation]

This operation loads a system.

The default methods for `load-op` compile files before loading them. For parity, your own methods on new component types should probably do so too.

`load-source-op`   [Operation]

This operation will load the source for the files in a module even if the source files have been compiled. Systems sometimes have knotty dependencies which require that sources are loaded before they can be compiled. This is how you do that.

If you are creating a component type, you need to implement this operation — at least, where meaningful.

`test-op`   [Operation]

This operation will perform some tests on the module. The default method will do nothing. The default dependency is to require `load-op` to be performed on the module first. The default `operation-done-p` is that the operation is *never* done — we assume that if you invoke the `test-op`, you want to test the system, even if you have already done so.

The results of this operation are not defined by ASDF. It has proven difficult to define how the test operation should signal its results to the user in a way that is compatible with all of the various test libraries and test techniques in use in the community.

6.1.2 Creating new operations

ASDF was designed to be extensible in an object-oriented fashion. To teach ASDF new tricks, a programmer can implement the behaviour he wants by creating a subclass of `operation`.

ASDF's pre-defined operations are in no way "privileged", but it is requested that developers never use the `asdf` package for operations they develop themselves. The rationale for this rule is that we don't want to establish a "global asdf operation name registry", but also want to avoid name clashes.

An operation must provide methods for the following generic functions when invoked with an object of type `source-file`: *FIXME describe this better*

- `output-files` The `output-files` method determines where the method will put its files. It returns two values, a list of pathnames, and a boolean. If the boolean is `T` then the pathnames are marked not be translated by enclosing `:around` methods. If the boolean is `NIL` then enclosing `:around` methods may translate these pathnames, e.g. to ensure object files are somehow stored in some implementation-dependent cache.

- `perform` The `perform` method must call `output-files` to find out where to put its files, because the user is allowed to override.

- `output-files` for local policy `explain`

- `operation-done-p`, if you don't like the default one

Operations that print output should send that output to the standard CL stream `*standard-output*`, as the Lisp compiler and loader do.

6.2 Components

A *component* represents a source file or (recursively) a collection of components. A *system* is (roughly speaking) a top-level component that can be found via `find-system`.

A *system designator* is a string or symbol and behaves just like any other component name (including with regard to the case conversion rules for component names).

`find-system` *system-designator* **&optional** (*error-p* **t**)   [Function]

> Given a system designator, `find-system` finds and returns a system. If no system is found, an error of type `missing-component` is thrown, or `nil` is returned if `error-p` is false.
>
> To find and update systems, `find-system` funcalls each element in the `*system-definition-search-functions*` list, expecting a pathname to be returned. The resulting pathname is loaded if either of the following conditions is true:
>
> - there is no system of that name in memory
> - the file's `last-modified` time exceeds the `last-modified` time of the system in memory
>
> When system definitions are loaded from '`.asd`' files, a new scratch package is created for them to load into, so that different systems do not overwrite each others operations. The user may also wish to (and is recommended to) include `defpackage` and `in-package` forms in his system definition files, however, so that they can be loaded manually if need be.
>
> The default value of `*system-definition-search-functions*` is a list of two functions. The first function looks in each of the directories given by evaluating members of `*central-registry*` for a file whose name is the name of the system and whose type is '`asd`'. The first such file is returned, whether or not it turns out to actually define the appropriate system. The second function does something similar, for the directories specified in the `source-registry`. Hence, it is strongly advised to define a system *foo* in the corresponding file *foo.asd*.

6.2.1 Common attributes of components

All components, regardless of type, have the following attributes. All attributes except `name` are optional.

6.2.1.1 Name

A component name is a string or a symbol. If a symbol, its name is taken and lowercased.

Unless overridden by a `:pathname` attribute, the name will be interpreted as a pathname specifier according to a Unix-style syntax. See Section 5.3 [Pathname specifiers], page 9.

6.2.1.2 Version identifier

This optional attribute is used by the `test-system-version` operation. See Section 6.1.1 [Predefined operations of ASDF], page 14. For the default method of `test-system-version`, the version should be a string of integers separated by dots, for example '`1.0.11`'.

Nota Bene: This operation, planned for ASDF 1, is still not implement yet as of ASDF 2. Don't hold your breath.

6.2.1.3 Required features

FIXME: This subsection seems to contradict the `defsystem` grammar subsection, which doesn't provide any obvious way to specify required features. Furthermore, in 2009, discussions on the asdf-devel mailing list suggested that the specification of required features may be broken, and that no one may have been using them for a while. Please contact the asdf-devel mailing list if you are interested in getting this features feature fixed.

Traditionally defsystem users have used reader conditionals to include or exclude specific per-implementation files. This means that any single implementation cannot read the entire system, which becomes a problem if it doesn't wish to compile it, but instead for example to create an archive file containing all the sources, as it will omit to process the system-dependent sources for other systems.

Each component in an asdf system may therefore specify features using the same syntax as `#+` does, and it will (somehow) be ignored for certain operations unless the feature conditional is a member of `*features*`.

6.2.1.4 Dependencies

This attribute specifies dependencies of the component on its siblings. It is optional but often necessary.

There is an excitingly complicated relationship between the initarg and the method that you use to ask about dependencies

Dependencies are between (operation component) pairs. In your initargs for the component, you can say

```
:in-order-to ((compile-op (load-op "a" "b") (compile-op "c"))
              (load-op (load-op "foo")))
```

This means the following things:

- before performing compile-op on this component, we must perform load-op on *a* and *b*, and compile-op on *c*,
- before performing `load-op`, we have to load *foo*

The syntax is approximately

```
(this-op {(other-op required-components)}+)

required-components := component-name
                     | (required-components required-components)

component-name := string
                | (:version string minimum-version-object)
```

Side note:

This is on a par with what ACL defsystem does. mk-defsystem is less general: it has an implied dependency

```
  for all x, (load x) depends on (compile x)
```

and using a `:depends-on` argument to say that *b* depends on *a* *actually* means that

```
(compile b) depends on (load a)
```

This is insufficient for e.g. the McCLIM system, which requires that all the files are loaded before any of them can be compiled]

End side note

In ASDF, the dependency information for a given component and operation can be queried using `(component-depends-on operation component)`, which returns a list

```
((load-op "a") (load-op "b") (compile-op "c") ...)
```

`component-depends-on` can be subclassed for more specific component/operation types: these need to `(call-next-method)` and append the answer to their dependency, unless they have a good reason for completely overriding the default dependencies.

If it weren't for CLISP, we'd be using `LIST` method combination to do this transparently. But, we need to support CLISP. If you have the time for some CLISP hacking, I'm sure they'd welcome your fixes.

6.2.1.5 pathname

This attribute is optional and if absent (which is the usual case), the component name will be used.

See Section 5.3 [Pathname specifiers], page 9, for an explanation of how this attribute is interpreted.

Note that the `defsystem` macro (used to create a "top-level" system) does additional processing to set the filesystem location of the top component in that system. This is detailed elsewhere. See Chapter 5 [Defining systems with defsystem], page 8.

The answer to the frequently asked question "how do I create a system definition where all the source files have a '`.cl`' extension" is thus

```
(defmethod source-file-type ((c cl-source-file) (s (eql (find-system 'my-sys))))
  "cl")
```

6.2.1.6 properties

This attribute is optional.

Packaging systems often require information about files or systems in addition to that specified by ASDF's pre-defined component attributes. Programs that create vendor packages out of ASDF systems therefore have to create "placeholder" information to satisfy these systems. Sometimes the creator of an ASDF system may know the additional information and wish to provide it directly.

`(component-property component property-name)` and associated `setf` method will allow the programmatic update of this information. Property names are compared as if by `EQL`, so use symbols or keywords or something.

6.2.2 Pre-defined subclasses of component

source-file   [Component]

>   A source file is any file that the system does not know how to generate from other components of the system.

>   Note that this is not necessarily the same thing as "a file containing data that is typically fed to a compiler". If a file is generated by some pre-processor stage (e.g.

a '`.h`' file from '`.h.in`' by autoconf) then it is not, by this definition, a source file. Conversely, we might have a graphic file that cannot be automatically regenerated, or a proprietary shared library that we received as a binary: these do count as source files for our purposes.

Subclasses of source-file exist for various languages. *FIXME: describe these.*

`module`   [Component]

A module is a collection of sub-components.

A module component has the following extra initargs:

- `:components` the components contained in this module
- `:default-component-class` All children components which don't specify their class explicitly are inferred to be of this type.
- `:if-component-dep-fails` This attribute takes one of the values `:fail`, `:try-next`, `:ignore`, its default value is `:fail`. The other values can be used for implementing conditional compilation based on implementation `*features*`, for the case where it is not necessary for all files in a module to be compiled. *FIXME: such conditional compilation has been reported to be broken in 2009.*
- `:serial` When this attribute is set, each subcomponent of this component is assumed to depend on all subcomponents before it in the list given to `:components`, i.e. all of them are loaded before a compile or load operation is performed on it.

The default operation knows how to traverse a module, so most operations will not need to provide methods specialised on modules.

`module` may be subclassed to represent components such as foreign-language linked libraries or archive files.

`system`   [Component]

`system` is a subclass of `module`.

A system is a module with a few extra attributes for documentation purposes; these are given elsewhere. See Section 5.3 [The defsystem grammar], page 9.

Users can create new classes for their systems: the default `defsystem` macro takes a `:class` keyword argument.

6.2.3 Creating new component types

New component types are defined by subclassing one of the existing component classes and specializing methods on the new component class.

FIXME: this should perhaps be explained more throughly, not only by example …

As an example, suppose we have some implementation-dependent functionality that we want to isolate in one subdirectory per Lisp implementation our system supports. We create a subclass of `cl-source-file`:

```
(defclass unportable-cl-source-file (cl-source-file)
  ())
```

A hypothetical function `system-dependent-dirname` gives us the name of the subdirectory. All that's left is to define how to calculate the pathname of an `unportable-cl-source-file`.

```
(defmethod component-pathname ((component unportable-cl-source-file))
  (let ((pathname (call-next-method))
        (name (string-downcase (system-dependent-dirname))))
    (merge-pathnames*
     (make-pathname :directory (list :relative name))
     pathname)))
```

The new component type is used in a `defsystem` form in this way:

```
(defsystem :foo
   :components
   ((:file "packages")
    ...
    (:unportable-cl-source-file "threads"
     :depends-on ("packages" ...))
    ...
    )
```

7 Controlling where ASDF searches for systems

7.1 Configurations

Configurations specify paths where to find system files.

1. The search registry may use some hardcoded wrapping registry specification. This allows some implementations (notably SBCL) to specify where to find some special implementation-provided systems that need to precisely match the version of the implementation itself.

2. An application may explicitly initialize the source-registry configuration using the configuration API (see Chapter 7 [Configuration API], page 21, below) in which case this takes precedence. It may itself compute this configuration from the command-line, from a script, from its own configuration file, etc.

3. The source registry will be configured from the environment variable `CL_SOURCE_REGISTRY` if it exists.

4. The source registry will be configured from user configuration file `$XDG_CONFIG_DIRS/common-lisp/source-registry.conf` (which defaults to `~/.config/common-lisp/source-registry.conf`) if it exists.

5. The source registry will be configured from user configuration directory `$XDG_CONFIG_DIRS/common-lisp/source-registry.conf.d/` (which defaults to `~/.config/common-lisp/source-registry.conf.d/`) if it exists.

6. The source registry will be configured from system configuration file `/etc/common-lisp/source-registry.conf` if it exists/

7. The source registry will be configured from system configuration directory `/etc/common-lisp/source-registry.conf.d/` if it exists.

8. The source registry will be configured from a default configuration. This configuration may allow for implementation-specific systems to be found, for systems to be found the current directory (at the time that the configuration is initialized) as well as `:directory` entries for `$XDG_DATA_DIRS/common-lisp/systems/` and `:tree` entries for `$XDG_DATA_DIRS/common-lisp/source/`.

Each of these configuration is specified as a SEXP in a trival domain-specific language (defined below). Additionally, a more shell-friendly syntax is available for the environment variable (defined yet below).

Each of these configurations is only used if the previous configuration explicitly or implicitly specifies that it includes its inherited configuration.

Additionally, some implementation-specific directories may be automatically prepended to whatever directories are specified in configuration files, no matter if the last one inherits or not.

7.2 XDG base directory

Note that we purport to respect the XDG base directory specification as to where configuration files are located, where data files are located, where output file caches are located. Mentions of XDG variables refer to that document.

http://standards.freedesktop.org/basedir-spec/basedir-spec-latest.html

This specification allows the user to specify some environment variables to customize how applications behave to his preferences.

On Windows platforms, when not using Cygwin, instead of the XDG base directory specification, we try to use folder configuration from the registry regarding `Common AppData` and similar directories. However, support querying the Windows registry is limited as of ASDF 2, and on many implementations, we may fall back to always using the defaults without consulting the registry. Patches welcome.

7.3 Backward Compatibility

For backward compatibility as well as for a practical backdoor for hackers, ASDF will first search for `.asd` files in the directories specified in `asdf:*central-registry*` before it searches in the source registry above.

See Chapter 3 [Configuring ASDF to find your systems – old style], page 4.

By default, `asdf:*central-registry*` will be empty.

This old mechanism will therefore not affect you if you don't use it, but will take precedence over the new mechanism if you do use it.

7.4 Configuration DSL

Here is the grammar of the SEXP DSL for source-registry configuration:

```
;; A configuration is single SEXP starting with keyword :source-registry
;; followed by a list of directives.
CONFIGURATION := (:source-registry DIRECTIVE ...)

;; A directive is one of the following:
DIRECTIVE :=
    ;; INHERITANCE DIRECTIVE:
    ;; Your configuration expression MUST contain
    ;; exactly one of either of these:
    :inherit-configuration | ; splices inherited configuration (often specified last)
    :ignore-inherited-configuration | ; drop inherited configuration (specified anywhe

    ;; add a single directory to be scanned (no recursion)
    (:directory DIRECTORY-PATHNAME-DESIGNATOR) |

    ;; add a directory hierarchy, recursing but excluding specified patterns
    (:tree DIRECTORY-PATHNAME-DESIGNATOR) |

    ;; override the defaults for exclusion patterns
    (:exclude PATTERN ...) |
    ;; augment the defaults for exclusion patterns
    (:also-exclude PATTERN ...) |

    ;; splice the parsed contents of another config file
```

```
        (:include REGULAR-FILE-PATHNAME-DESIGNATOR) |

        ;; This directive specifies that some default must be spliced.
        :default-registry

    PATTERN := a string without wildcards, that will be matched exactly
    against the name of a any subdirectory in the directory component
            of a path. e.g. "_darcs" will match '#p"/foo/bar/_darcs/src/bar.asd"'▮
```

For instance, as a simple case, my '`~/.config/common-lisp/source-registry.conf`', which is the default place ASDF looks for this configuration, once contained:

```
(:source-registry
  (:tree "/home/fare/cl/")
  :inherit-configuration)
```

7.5 Configuration Directories

Configuration directories consist in files each contains a list of directives without any enclosing (`:source-registry ...`) form. The files will be sorted by namestring as if by `string<` and the lists of directives of these files with be concatenated in order. An implicit `:inherit-configuration` will be included at the *end* of the list.

This allows for packaging software that has file granularity (e.g. Debian's `dpkg` or some future version of `clbuild`) to easily include configuration information about distributed software.

The convention is that, for sorting purposes, the names of files in such a directory begin with two digits that determine the order in which these entries will be read. Also, the type of these files is conventionally `"conf"` and as a limitation to some implementations (e.g. GNU clisp), the type cannot be `NIL`.

Directories may be included by specifying a directory pathname or namestring in an `:include` directive, e.g.:

```
    (:include "/foo/bar/")
```

Hence, to achieve the same effect as my example '`~/.config/common-lisp/source-registry.conf`'▮ above, I could simply create a file '`~/.config/common-lisp/source-registry.conf.d/33-home-fare-cl.con`' alone in its directory with the following contents:

```
    (:tree "/home/fare/cl/")
```

7.6 Shell-friendly syntax for configuration

When considering environment variable `CL_SOURCE_REGISTRY` ASDF will skip to next configuration if it's an empty string. It will `READ` the string as a SEXP in the DSL if it begins with a paren `(` and it will be interpreted much like `TEXINPUTS` list of paths, where

* paths are separated by a `:` (colon) on Unix platforms (including cygwin), by a `;` (semicolon) on other platforms (mainly, Windows).

* each entry is a directory to add to the search path.

* if the entry ends with a double slash `//` then it instead indicates a tree in the subdirectories of which to recurse.

* if the entry is the empty string (which may only appear once), then it indicates that the inherited configuration should be spliced there.

7.7 Search Algorithm

In case that isn't clear, the semantics of the configuration is that when searching for a system of a given name, directives are processed in order.

When looking in a directory, if the system is found, the search succeeds, otherwise it continues.

When looking in a tree, if one system is found, the search succeeds. If multiple systems are found, the consequences are unspecified: the search may succeed with any of the found systems, or an error may be raised. ASDF currently returns the first system found, XCVB currently raised an error. If none is found, the search continues.

Exclude statements specify patterns of subdirectories the systems from which to ignore. Typically you don't want to use copies of files kept by such version control systems as Darcs. Exclude statements are not propagated to further included or inherited configuration files or expressions; instead the defaults are reset around every configuration statement to the default defaults from `asdf::*default-source-registry-exclusions*`.

Include statements cause the search to recurse with the path specifications from the file specified.

An inherit-configuration statement cause the search to recurse with the path specifications from the next configuration (see Chapter 7 [Configurations], page 21 above).

7.8 Caching Results

The implementation is allowed to either eagerly compute the information from the configurations and file system, or to lazily re-compute it every time, or to cache any part of it as it goes. To explicitly flush any information cached by the system, use the API below.

7.9 Configuration API

The specified functions are exported from your build system's package. Thus for ASDF the corresponding functions are in package ASDF, and for XCVB the corresponding functions are in package XCVB.

`initialize-source-registry` *&optional PARAMETER*   [Function]
    will read the configuration and initialize all internal variables. You may extend or override configuration from the environment and configuration files with the given *PARAMETER*, which can be `NIL` (no configuration override), or a SEXP (in the SEXP DSL), a string (as in the string DSL), a pathname (of a file or directory with configuration), or a symbol (fbound to function that when called returns one of the above).

`clear-source-registry`   [Function]
    undoes any source registry configuration and clears any cache for the search algorithm. You might want to call that before you dump an image that would be resumed with a different configuration, and return an empty configuration. Note that this does not include clearing information about systems defined in the current image, only about where to look for systems not yet defined.

`ensure-source-registry` *&optional PARAMETER*   [Function]
>   checks whether a source registry has been initialized. If not, initialize it with the
>   given *PARAMETER*.

7.10 Future

If this mechanism is successful, in the future, we may declare `asdf:*central-registry*`
obsolete and eventually remove it. Any hook into implementation-specific search mecha-
nisms will by then have been integrated in the `:default-configuration` which everyone
should either explicitly use or implicit inherit. Some shell syntax for it should probably be
added somehow.

But we're not there yet. For now, let's see how practical this new source-registry is.

7.11 Rejected ideas

Alternatives I considered and rejected included:

1. Keep `asdf:*central-registry*` as the master with its current semantics, and some-
   how the configuration parser expands the new configuration language into a expanded
   series of directories of subdirectories to lookup, pre-recursing through specified hierar-
   chies. This is kludgy, and leaves little space of future cleanups and extensions.

2. Keep `asdf:*central-registry*` remains the master but extend its semantics in com-
   pletely new ways, so that new kinds of entries may be implemented as a recursive
   search, etc. This seems somewhat backwards.

3. Completely remove `asdf:*central-registry*` and break backwards compatibility.
   Hopefully this will happen in a few years after everyone migrate to a better ASDF
   and/or to XCVB, but it would be very bad to do it now.

4. Replace `asdf:*central-registry*` by a symbol-macro with appropriate magic when
   you dereference it or setf it. Only the new variable with new semantics is handled by
   the new search procedure. Complex and still introduces subtle semantic issues.

I've been suggested the below features, but have rejected them, for the sake of keeping
ASDF no more complex than strictly necessary.

- More syntactic sugar: synonyms for the configuration directives, such as
  `(:add-directory X)` for `(:directory X)`, or `(:add-directory-hierarchy X)` or
  `(:add-directory X :recurse t)` for `(:tree X)`.

- The possibility to register individual files instead of directories.

- Integrate Xach Beane's tilde expander into the parser, or something similar that is shell-
  friendly or shell-compatible. I'd rather keep ASDF minimal. But maybe this precisely
  keeps it minimal by removing the need for evaluated entries that ASDF has? i.e. uses
  of `USER-HOMEDIR-PATHNAME` and `$SBCL_HOME` Hopefully, these are already superseded
  by the `:default-registry`

- Using the shell-unfriendly syntax `/**` instead of `//` to specify recursion down a filesys-
  tem tree in the environment variable. It isn't that Lisp friendly either.

7.12 TODO

- Add examples

7.13 Credits for the source-registry

Thanks a lot to Stelian Ionescu for the initial idea.

Thanks to Rommel Martinez for the initial implementation attempt.

All bad design ideas and implementation bugs are to mine, not theirs. But so are good design ideas and elegant implementation tricks.

— Francois-Rene Rideau `fare@tunes.org`, Mon, 22 Feb 2010 00:07:33 -0500

8 Controlling where ASDF saves compiled files

Each Common Lisp implementation has its own format for compiled files (fasls for short, short for "fast loading"). If you use multiple implementations (or multiple versions of the same implementation), you'll soon find your source directories littered with various 'fasl's, 'dfsl's, 'cfsl's and so on. Worse yet, some implementations use the same file extension while changing formats from version to version (or platform to platform) which means that you'll have to recompile binaries as you switch from one implementation to the next.

ASDF 2 includes the `asdf-output-translations` facility to mitigate the problem.

8.1 Configurations

Configurations specify mappings from input locations to output locations. Once again we rely on the XDG base directory specification for configuration. See Chapter 7 [XDG base directory], page 21.

 1. Some hardcoded wrapping output translations configuration may be used. This allows special output translations (or usually, invariant directories) to be specified corresponding to the similar special entries in the source registry.

 2. An application may explicitly initialize the output-translations configuration using the Configuration API in which case this takes precedence. (see Chapter 8 [Configuration API], page 27.) It may itself compute this configuration from the command-line, from a script, from its own configuration file, etc.

 3. The source registry will be configured from the environment variable `ASDF_OUTPUT_TRANSLATIONS` if it exists.

 4. The source registry will be configured from user configuration file '`$XDG_CONFIG_DIRS/common-lisp/asdf-output-translations.conf`' (which defaults to '`~/.config/common-lisp/asdf-output-translations.conf`') if it exists.

 5. The source registry will be configured from user configuration directory '`$XDG_CONFIG_DIRS/common-lisp/asdf-output-translations.conf.d/`' (which defaults to '`~/.config/common-lisp/asdf-output-translations.conf.d/`') if it exists.

 6. The source registry will be configured from system configuration file '`/etc/common-lisp/asdf-output-translations.conf`' if it exists.

 7. The source registry will be configured from system configuration directory '`/etc/common-lisp/asdf-output-translations.conf.d/`' if it exists.

Each of these configurations is specified as a SEXP in a trival domain-specific language (defined below). Additionally, a more shell-friendly syntax is available for the environment variable (defined yet below).

Each of these configurations is only used if the previous configuration explicitly or implicitly specifies that it includes its inherited configuration.

Note that by default, a per-user cache is used for output files. This allows the seamless use of shared installations of software between several users, and takes files out of the way of the developers when they browse source code, at the expense of taking a small toll when developers have to clean up output files and find they need to get familiar with output-translations first.

8.2 Backward Compatibility

We purposefully do NOT provide backward compatibility with earlier versions of `ASDF-Binary-Locations` (8 Sept 2009), `common-lisp-controller` (7.0) or `cl-launch` (2.35), each of which had similar general capabilities. The previous APIs of these programs were not designed for configuration by the end-user in an easy way with configuration files. Recent versions of same packages use the new `asdf-output-translations` API as defined below: `common-lisp-controller` (7.2) and `cl-launch` (3.000). `ASDF-Binary-Locations` is fully superseded and not to be used anymore.

This incompatibility shouldn't inconvenience many people. Indeed, few people use and customize these packages; these few people are experts who can trivially adapt to the new configuration. Most people are not experts, could not properly configure these features (except inasmuch as the default configuration of `common-lisp-controller` and/or `cl-launch` might have been doing the right thing for some users), and yet will experience software that "just works", as configured by the system distributor, or by default.

Nevertheless, if you are a fan of `ASDF-Binary-Locations`, we provide a limited emulation mode:

`asdf:enable-asdf-binary-locations-compatibility` *&key*   [Function]
> *centralize-lisp-binaries default-toplevel-directory include-per-user-information*
> *map-all-source-files source-to-target-mappings*
>> This function will initialize the new `asdf-output-translations` facility in a way that emulates the behavior of the old `ASDF-Binary-Locations` facility. Where you would previously set global variables *centralize-lisp-binaries*, *default-toplevel-directory*, *include-per-user-information*, *map-all-source-files* or *source-to-target-mappings* you will now have to pass the same values as keyword arguments to this function. Note however that as an extension the `:source-to-target-mappings` keyword argument will accept any valid pathname designator for `asdf-output-translations` instead of just strings and pathnames.

If you insist, you can also keep using the old `ASDF-Binary-Locations` (the one available as an extension to load of top of ASDF, not the one built into a few old versions of ASDF), but first you must disable `asdf-output-translations` with (`asdf:disable-output-translations`), or you might experience "interesting" issues.

Also, note that output translation is enabled by default. To disable it, use (`asdf:disable-output-translations`).

8.3 Configuration DSL

Here is the grammar of the SEXP DSL for `asdf-output-translations` configuration:

```
;; A configuration is single SEXP starting with keyword :source-registry
;; followed by a list of directives.
CONFIGURATION := (:output-translations DIRECTIVE ...)

;; A directive is one of the following:
DIRECTIVE :=
    ;; INHERITANCE DIRECTIVE:
    ;; Your configuration expression MUST contain
```

```
    ;; exactly one of either of these:
    :inherit-configuration | ; splices inherited configuration (often specified last)█
    :ignore-inherited-configuration | ; drop inherited configuration (specified anywhere)█

    ;; include a configuration file or directory
    (:include PATHNAME-DESIGNATOR) |

    ;; enable global cache in ~/.common-lisp/cache/sbcl-1.0.35-x86-64/ or something.█
    :enable-user-cache |
    ;; Disable global cache. Map / to /
    :disable-cache |

    ;; add a single directory to be scanned (no recursion)
    (DIRECTORY-DESIGNATOR DIRECTORY-DESIGNATOR)

    ;; use a function to return the translation of a directory designator
    (DIRECTORY-DESIGNATOR (:function TRANSLATION-FUNCTION))

DIRECTORY-DESIGNATOR :=
    T | ;; as source matches anything, as destination leaves pathname unmapped.█
    ABSOLUTE-COMPONENT-DESIGNATOR |
    (ABSOLUTE-COMPONENT-DESIGNATOR RELATIVE-COMPONENT-DESIGNATOR ...)

ABSOLUTE-COMPONENT-DESIGNATOR :=
    NULL | ;; As source: skip this entry. As destination: same as source
    :ROOT | ;; magic: paths that are relative to the root of the source host and device█
    STRING | ;; namestring (directory is assumed, better be absolute or bust, ''/**/*.*'' a
    PATHNAME | ;; pathname (better be an absolute directory or bust)
    :HOME | ;; designates the user-homedir-pathname ~/
    :USER-CACHE | ;; designates the default location for the user cache
    :SYSTEM-CACHE | ;; designates the default location for the system cache
    :CURRENT-DIRECTORY ;; the current directory

RELATIVE-COMPONENT-DESIGNATOR :=
    STRING | ;; namestring, directory is assumed. If the last component, /**/*.* is added█
    PATHNAME | ;; pathname unless last component, directory is assumed.
    :IMPLEMENTATION | ;; a directory based on implementation, e.g. sbcl-1.0.32.30-linux-x86
    :IMPLEMENTATION-TYPE | ;; a directory based on lisp-implementation-type only, e.g. sbcl
    :CURRENT-DIRECTORY | ;; all components of the current directory, without the :absolute█
    :UID | ;; current UID -- not available on Windows
    :USER ;; current USER name -- NOT IMPLEMENTED(!)

TRANSLATION-FUNCTION :=
    SYMBOL | ;; symbol of a function that takes two arguments,
             ;; the pathname to be translated and the matching DIRECTORY-DESIGNATOR█
    LAMBDA   ;; A form which evalutates to a function taking two arguments consisting of█
             ;; the pathname to be translated and the matching DIRECTORY-DESIGNATOR█
```

Relative components better be either relative or subdirectories of the path before them, or bust.

The last component, if not a pathname, is notionally completed by '`/**/*.*`'. You can specify more fine-grained patterns by using a pathname object as the last component e.g. '`#p"some/path/**/foo*/bar-*.fasl"`'

You may use `#+features` to customize the configuration file.

The second designator of a mapping may be `NIL`, indicating that files are not mapped to anything but themselves (same as if the second designator was the same as the first).

When the first designator is `t`, the mapping always matches. When the first designator starts with `:root`, the mapping matches any host and device. In either of these cases, if the second designator isn't `t` and doesn't start with `:root`, then strings indicating the host and pathname are somehow copied in the beginning of the directory component of the source pathname before it is translated.

When the second designator is `t`, the mapping is the identity. When the second designator starts with `root`, the mapping preserves the host and device of the original pathname.

`:include` statements cause the search to recurse with the path specifications from the file specified.

If the `translate-pathname` mechanism cannot achieve a desired translation, the user may provide a function which provides the required algorithim. Such a translation function is specified by supplying a list as the second `directory-designator` the first element of which is the keyword `:function`, and the second element of which is either a symbol which designates a function or a lambda expression. The function designated by the second argument must take two arguments, the first being the pathname of the source file, the second being the wildcard that was matched. The result of the function invocation should be the translated pathname.

An `:inherit-configuration` statement cause the search to recurse with the path specifications from the next configuration. See Chapter 8 [Configurations], page 27, above.

- `:enable-user-cache` is the same as (`t :user-cache`).
- `:disable-cache` is the same as (`t t`).
- `:user-cache` uses the contents of variable `asdf::*user-cache*` which by default is the same as using (`:home ".cache" "common-lisp" :implementation`).
- `:system-cache` uses the contents of variable `asdf::*system-cache*` which by default is the same as using (`"/var/cache/common-lisp" :uid :implementation-type`) (on Unix and cygwin), or something semi-sensible on Windows.

8.4 Configuration Directories

Configuration directories consist in files each contains a list of directives without any enclosing (`:output-translations ...`) form. The files will be sorted by namestring as if by `string<` and the lists of directives of these files with be concatenated in order. An implicit `:inherit-configuration` will be included at the *end* of the list.

This allows for packaging software that has file granularity (e.g. Debian's `dpkg` or some future version of `clbuild`) to easily include configuration information about software being distributed.

The convention is that, for sorting purposes, the names of files in such a directory begin with two digits that determine the order in which these entries will be read. Also, the type of these files is conventionally `"conf"` and as a limitation of some implementations, the type cannot be `NIL`.

Directories may be included by specifying a directory pathname or namestring in an `:include` directive, e.g.:

```
(:include "/foo/bar/")
```

8.5 Shell-friendly syntax for configuration

When considering environment variable `ASDF_OUTPUT_TRANSLATIONS` ASDF will skip to next configuration if it's an empty string. It will `READ` the string as an SEXP in the DSL if it begins with a paren `(` and it will be interpreted as a list of directories. Directories should come by pairs, indicating a mapping directive. Entries are separated by a `:` (colon) on Unix platforms (including cygwin), by a `;` (semicolon) on other platforms (mainly, Windows).

The magic empty entry, if it comes in what would otherwise be the first entry in a pair, indicates the splicing of inherited configuration. If it comes as the second entry in a pair, it indicates that the directory specified first is to be left untranslated (which has the same effect as if the directory had been repeated).

8.6 Semantics of Output Translations

From the specified configuration, a list of mappings is extracted in a straightforward way: mappings are collected in order, recursing through included or inherited configuration as specified. To this list is prepended some implementation-specific mappings, and is appended a global default.

The list is then compiled to a mapping table as follows: for each entry, in order, resolve the first designated directory into an actual directory pathname for source locations. If no mapping was specified yet for that location, resolve the second designated directory to an output location directory add a mapping to the table mapping the source location to the output location, and add another mapping from the output location to itself (unless a mapping already exists for the output location).

Based on the table, a mapping function is defined, mapping source pathnames to output pathnames: given a source pathname, locate the longest matching prefix in the source column of the mapping table. Replace that prefix by the corresponding output column in the same row of the table, and return the result. If no match is found, return the source pathname. (A global default mapping the filesystem root to itself may ensure that there will always be a match, with same fall-through semantics).

8.7 Caching Results

The implementation is allowed to either eagerly compute the information from the configurations and file system, or to lazily re-compute it every time, or to cache any part of it as it goes. To explicitly flush any information cached by the system, use the API below.

8.8 Output location API

The specified functions are exported from package ASDF.

`initialize-output-translations` *&optional PARAMETER*   [Function]

> will read the configuration and initialize all internal variables. You may extend or override configuration from the environment and configuration files with the given *PARAMETER*, which can be `NIL` (no configuration override), or a SEXP (in the SEXP DSL), a string (as in the string DSL), a pathname (of a file or directory with configuration), or a symbol (fbound to function that when called returns one of the above).

`disable-output-translations`   [Function]

> will initialize output translations in a way that maps every pathname to itself, effectively disabling the output translation facility.

`clear-output-translations`   [Function]

> undoes any output translation configuration and clears any cache for the mapping algorithm. You might want to call that before you dump an image that would be resumed with a different configuration, and return an empty configuration. Note that this does not include clearing information about systems defined in the current image, only about where to look for systems not yet defined.

`ensure-output-translations` *&optional PARAMETER*   [Function]

> checks whether output translations have been initialized. If not, initialize them with the given *PARAMETER*. This function will be called before any attempt to operate on a system.

`apply-output-translations` *PATHNAME*   [Function]

> Applies the configured output location translations to *PATHNAME* (calls `ensure-output-translations` for the translations).

8.9 Credits for output translations

Thanks a lot to Bjorn Lindberg and Gary King for `ASDF-Binary-Locations`, and to Peter van Eynde for `Common Lisp Controller`.

All bad design ideas and implementation bugs are to mine, not theirs. But so are good design ideas and elegant implementation tricks.

— Francois-Rene Rideau fare@tunes.org

9 Error handling

9.1 ASDF errors

If ASDF detects an incorrect system definition, it will signal a generalised instance of
`SYSTEM-DEFINITION-ERROR`.

Operations may go wrong (for example when source files contain errors). These are
signalled using generalised instances of `OPERATION-ERROR`.

9.2 Compilation error and warning handling

ASDF checks for warnings and errors when a file is compiled. The variables *compile-file-warnings-behaviour* and *compile-file-errors-behavior* control the handling of any such
events. The valid values for these variables are `:error`, `:warn`, and `:ignore`.

10 Miscellaneous additional functionality

FIXME: Add discussion of `run-shell-command`*? Others?*

ASDF includes several additional features that are generally useful for system definition and development. These include:

system-relative-pathname *system name &key type*   [Function]

It's often handy to locate a file relative to some system. The `system-relative-pathname` function meets this need. It takes two arguments: the name of a system and a relative pathname. It returns a pathname built from the location of the system's source file and the relative pathname. For example

```
> (asdf:system-relative-pathname 'cl-ppcre #p"regex.data")
#P"/repository/other/cl-ppcre/regex.data"
```

Instead of a pathname, you can provide a symbol or a string, and optionally a keyword argument `type`. The arguments will then be interpreted in the same way as pathname specifiers for components. See .

system-source-directory *system-designator*   [Function]

ASDF does not provide a turnkey solution for locating data (or other miscellaneous) files that are distributed together with the source code of a system. Programmers can use `system-source-directory` to find such files. Returns a pathname object. The *system-designator* may be a string, symbol, or ASDF system object.

clear-system *system-designator*   [Function]

It is sometimes useful to force recompilation of a previously loaded system. In these cases, it may be useful to (`asdf:clear-system :foo`) to remove the system from the table of currently loaded systems; the next time the system `foo` or one that depends on it is re-loaded, `foo` will then be loaded again. Alternatively, you could touch `foo.asd` or remove the corresponding fasls from the output file cache. (It was once conceived that one should provide a list of systems the recompilation of which to force as the `:force` keyword argument to `load-system`; but this has never worked, and though the feature was fixed in ASDF 2.000, it remains `cerror`'ed out as nobody ever used it.)

Note that this does not and cannot by itself undo the previous loading of the system. Common Lisp has no provision for such an operation, and its reliance on irreversible side-effects to global datastructures makes such a thing impossible in the general case. If the software being re-loaded is not conceived with hot upgrade in mind, this re-loading may cause many errors, warnings or subtle silent problems, as packages, generic function signatures, structures, types, macros, constants, etc. are being redefined incompatibly. It is up to the user to make sure that reloading is possible and has the desired effect. In some cases, extreme measures such as recursively deleting packages, unregistering symbols, defining methods on `update-instance-for-redefined-class` and much more are necessary for reloading to happen smoothly. ASDF itself goes through notable pains to make such a hot upgrade possible with respect to its own code, and what it does is ridiculously complex; look at the beginning of '`asdf.lisp`' to see what it does.

11 Getting the latest version

Decide which version you want. HEAD is the newest version and usually OK, whereas
RELEASE is for cautious people (e.g. who already have systems using ASDF that they
don't want broken), a slightly older version about which none of the HEAD users have
complained. There is also a STABLE version, which is earlier than release.

You may get the ASDF source repository using git: `git clone git://common-lisp.net/projects/asdf/asdf.git`

You will find the above referenced tags in this repository. You can also browse the
repository on `http://common-lisp.net/gitweb?p=projects/asdf/asdf.git`.

Discussion of ASDF development is conducted on the mailing list `asdf-devel@common-lisp.net`. `http://common-lisp.net/cgi-bin/mailman/listinfo/asdf-devel`

12 FAQ

12.1 "Where do I report a bug?"

ASDF bugs are tracked on launchpad: https://launchpad.net/asdf.

If you're unsure about whether something is a bug, of for general discussion, use the asdf-devel mailing list

12.2 "What has changed between ASDF 1 and ASDF 2?"

12.2.1 What are ASDF 1 and ASDF 2?

On May 31st 2010, we have released ASDF 2. ASDF 2 refers to release 2.000 and later. (Releases between 1.656 and 1.728 were development releases for ASDF 2.) ASDF 1 to any release earlier than 1.369 or so. If your ASDF doesn't sport a version, it's an old ASDF 1.

ASDF 2 and its release candidates push `:asdf2` onto `*features*` so that if you are writing ASDF-dependent code you may check for this feature to see if the new API is present. *All* versions of ASDF should have the `:asdf` feature.

If you are experiencing problems or limitations of any sort with ASDF 1, we recommend that you should upgrade to ASDF 2, or whatever is the latest release.

12.2.2 ASDF can portably name files in subdirectories

Common Lisp namestrings are not portable, except maybe for logical pathnamestrings, that themselves have various limitations and require a lot of setup that is itself ultimately non-portable.

In ASDF 1, the only portable ways to refer to pathnames inside systems and components were very awkward, using `#.(make-pathname ...)` and `#.(merge-pathnames ...)`. Even the above were themselves were inadequate in the general case due to host and device issues, unless horribly complex patterns were used. Plenty of simple cases that looked portable actually weren't, leading to much confusion and greavance.

ASDF 2 implements its own portable syntax for strings as pathname specifiers. Naming files within a system definition becomes easy and portable again. See Chapter 10 [Miscellaneous additional functionality], page 34, `asdf-utilities:merge-pathnames*`, `asdf::merge-component-name-type`.

On the other hand, there are places where systems used to accept namestrings where you must now use an explicit pathname object: `(defsystem ... :pathname "LOGICAL-HOST:PATH;TO;SYSTEM;" ...)` must now be written with the `#p` syntax: `(defsystem ... :pathname #p"LOGICAL-HOST:PATH;TO;SYSTEM;" ...)`

See Section 5.3 [Pathname specifiers], page 9.

12.2.3 Output translations

A popular feature added to ASDF was output pathname translation: `asdf-binary-locations`, `common-lisp-controller`, `cl-launch` and other hacks were all implementing it in ways both mutually incompatible and difficult to configure.

Output pathname translation is essential to share source directories of portable systems across multiple implementations or variants thereof, or source directories of shared installations of systems across multiple users, or combinations of the above.

In ASDF 2, a standard mechanism is provided for that, `asdf-output-translations`, with sensible defaults, adequate configuration languages, a coherent set of configuration files and hooks, and support for non-Unix platforms.

See Chapter 8 [Controlling where ASDF saves compiled files], page 27.

12.2.4 Source Registry Configuration

Configuring ASDF used to require special magic to be applied just at the right moment, between the moment ASDF is loaded and the moment it is used, in a way that is specific to the user, the implementation he is using and the application he is building.

This made for awkward configuration files and startup scripts that could not be shared between users, managed by administrators or packaged by distributions.

ASDF 2 provides a well-documented way to configure ASDF, with sensible defaults, adequate configuration languages, and a coherent set of configuration files and hooks.

We believe it's a vast improvement because it decouples application distribution from library distribution. The application writer can avoid thinking where the libraries are, and the library distributor (dpkg, clbuild, advanced user, etc.) can configure them once and for every application. Yet settings can be easily overridden where needed, so whoever needs control has exactly as much as required.

At the same time, ASDF 2 remains compatible with the old magic you may have in your build scripts (using `*central-registry*` and `*system-definition-search-functions*`) to tailor the ASDF configuration to your build automation needs, and also allows for new magic, simpler and more powerful magic.

See Chapter 7 [Controlling where ASDF searches for systems], page 21.

12.2.5 Usual operations are made easier to the user

In ASDF 1, you had to use the awkward syntax (`asdf:oos 'asdf:load-op :foo`) to load a system, and similarly for `compile-op`, `test-op`.

In ASDF 2, you can use shortcuts for the usual operations: (`asdf:load-system :foo`), and similarly for `compile-system`, `test-system`.

12.2.6 Many bugs have been fixed

The following issues and many others have been fixed:

- The infamous TRAVERSE function has been revamped significantly, with many bugs squashed. In particular, dependencies were not correctly propagated across submodules within a system but now are. The :version and :feature features and the :force (system1 .. systemN) feature have been fixed.

- Performance has been notably improved for large systems (say with thousands of components) by using hash-tables instead of linear search, and linear-time list accumulation instead of quadratic-time recursive appends.

- Many features used to not be portable, especially where pathnames were involved. Windows support was notably quirky because of such non-portability.

- The internal test suite used to massively fail on many implementations. While still incomplete, it now fully passes on all implementations supported by the test suite, except for GCL (due to GCL bugs).

- Support was lacking for some implementations. ABCL and GCL were notably wholly broken. ECL extensions were not integrated in the ASDF release.

- The documentation was grossly out of date.

12.2.7 ASDF itself is versioned

Between new features, old bugs fixed, and new bugs introduced, there were various releases of ASDF in the wild, and no simple way to check which release had which feature set. People using or writing systems had to either make worst-case assumptions as to what features were available and worked, or take great pains to have the correct version of ASDF installed.

With ASDF 2, we provide a new stable set of working features that everyone can rely on from now on. Use `#+asdf2` to detect presence of ASDF 2, `(asdf:version-satisfies (asdf:asdf-version) "2.000")` to check the availability of a version no earlier than required.

12.2.8 ASDF can be upgraded

When an old version of ASDF was loaded, it was very hard to upgrade ASDF in your current image without breaking everything. Instead you have to exit the Lisp process and somehow arrange to start a new one from a simpler image. Something that can't be done from within Lisp, making automation of it difficult, which compounded with difficulty in configuration, made the task quite hard. Yet as we saw before, the task would have been required to not have to live with the worst case or non-portable subset of ASDF features.

With ASDF 2, it is easy to upgrade from ASDF 2 to later versions from within Lisp, and not too hard to upgrade from ASDF 1 to ASDF 2 from within Lisp. We support hot upgrade of ASDF and any breakage is a bug that we will do our best to fix. There are still limitations on upgrade, though, most notably the fact that after you upgrade ASDF, you must also reload or upgrade all ASDF extensions.

12.2.9 Decoupled release cycle

When vendors were releasing their Lisp implementations with ASDF, they had to basically never change version because neither upgrade nor downgrade was possible without breaking something for someone, and no obvious upgrade path was visible and recommendable.

With ASDF 2, upgrade is possible, easy and can be recommended. This means that vendors can safely ship a recent version of ASDF, confident that if a user isn't fully satisfied, he can easily upgrade ASDF and deal with a supported recent version of it. This means that release cycles will be causally decoupled, the practical consequence of which will mean faster convergence towards the latest version for everyone.

12.2.10 Pitfalls of ASDF 2

The main pitfalls in upgrading to ASDF 2 seem to be related to the output translation mechanism.

- Output translations is enabled by default. This may surprise some users, most of them in pleasant way (we hope), a few of them in an unpleasant way. It is trivial to disable output translations. See Chapter 12 ["How can I wholly disable the compiler output cache?"], page 36.

- Some systems in the large have been known not to play well with output translations. They were relatively easy to fix. Once again, it is also easy to disable output translations, or to override its configuration.

- The new ASDF output translations are incompatible with ASDF-Binary-Locations. They replace A-B-L, and there is compatibility mode to emulate your previous A-B-L configuration. See `asdf:enable-asdf-binary-locations-compatibility` in see Chapter 8 [Backward Compatibility], page 27. But thou shall not load ABL on top of ASDF 2.

- ASDF pathname designators are now specified in places where they were unspecified, and a few small adjustments have to be made to some non-portable defsystems. Notably, in the `:pathname` argument to a `defsystem` and its components, a logical pathname (or implementation-dependent hierarchical pathname) must now be specified with `#p` syntax where the namestring might have previously sufficed; moreover when evaluation is desired `#.` must be used, where it wasn't necessary in the toplevel `:pathname` argument.

Other issues include the following:

- There is a slight performance bug, notably on SBCL, when initially searching for 'asd' files, the implicit (`directory "/configured/path/**/*.asd"`) for every configured path (`:tree "/configured/path/"`) in your `source-registry` configuration can cause a slight pause. Try to (`time (asdf:initialize-source-registry)`) to see how bad it is or isn't on your system. If you insist on not having this pause, you can avoid the pause by overriding the default source-registry configuration and not use any deep `:tree` entry but only `:directory` entries or shallow `:tree` entries. Or you can fix your implementation to not be quite that slow when recursing through directories.

- On Windows, only LispWorks supports proper default configuration pathnames based on the Windows registry. Other implementations make do. Windows support is largely untested, so please help report and fix bugs.

12.3 Issues with installing the proper version of ASDF

12.3.1 "My Common Lisp implementation comes with an outdated version of ASDF. What to do?"

We recommend you upgrade ASDF. See Chapter 2 [Upgrading ASDF], page 2.

If this does not work, it is a bug, and you should report it. See Chapter 12 [Where do I report a bug], page 36. In the meantime, you can load 'asdf.lisp' directly. See Chapter 2 [Loading ASDF], page 2.

12.3.2 "I'm a Common Lisp implementation vendor. When and how should I upgrade ASDF?"

Starting with current candidate releases of ASDF 2, it should always be a good time to upgrade to a recent version of ASDF. You may consult with the maintainer for which

specific version they recommend, but the latest RELEASE should be correct. We trust you to thoroughly test it with your implementation before you release it. If there are any issues with the current release, it's a bug that you should report upstream and that we will fix ASAP.

As to how to include ASDF, we recommend the following:

- If ASDF isn't installed yet, then `(require :asdf)` should load the version of ASDF that is bundled with your system. You may have it load some other version configured by the user, if you allow such configuration.

- If your system provides a mechanism to hook into `CL:REQUIRE`, then it would be nice to add ASDF to this hook the same way that ABCL, CCL, CMUCL, ECL and SBCL do it.

- You may, like SBCL, have ASDF be implicitly used to require systems that are bundled with your Lisp distribution. If you do have a few magic systems that come with your implementation in a precompiled way such that one should only use the binary version that goes with your distribution, like SBCL does, then you should add them in the beginning of `wrapping-source-registry`.

- If you have magic systems as above, like SBCL does, then we explicitly ask you to *NOT* distribute 'asdf.asd' as part of those magic systems. You should still include the file 'asdf.lisp' in your source distribution and precompile it in your binary distribution, but 'asdf.asd' if included at all, should be secluded from the magic systems, in a separate file hierarchy, or you may otherwise rename the system and its file to e.g. `asdf-ecl` and 'asdf-ecl.asd', or `sb-asdf` and 'sb-asdf.asd'. Indeed, if you made 'asdf.asd' a magic system, then users would no longer be able to upgrade ASDF using ASDF itself to some version of their preference that they maintain independently from your Lisp distribution.

- If you do not have any such magic systems, or have other non-magic systems that you want to bundle with your implementation, then you may add them to the `default-source-registry`, and you are welcome to include 'asdf.asd' amongst them.

- Please send us upstream any patches you make to ASDF itself, so we can merge them back in for the benefit of your users when they upgrade to the upstream version.

12.4 Issues with configuring ASDF

12.4.1 "How can I customize where fasl files are stored?"

See Chapter 8 [Controlling where ASDF saves compiled files], page 27.

Note that in the past there was an add-on to ASDF called `ASDF-binary-locations`, developed by Gary King. That add-on has been merged into ASDF proper, then superseded by the `asdf-output-translations` facility.

Note that use of `asdf-output-translations` can interfere with one aspect of your systems — if your system uses `*load-truename*` to find files (e.g., if you have some data files stored with your program), then the relocation that this ASDF customization performs is likely to interfere. Use `asdf:system-relative-pathname` to locate a file in the source directory of some system, and use `asdf:apply-output-translations` to locate a file whose pathname has been translated by the facility.

12.4.2 "How can I wholly disable the compiler output cache?"

To permanently disable the compiler output cache for all future runs of ASDF, you can:

```
mkdir -p ~/.config/common-lisp/asdf-output-translations.conf.d/
echo ':disable-cache' > ~/.config/common-lisp/asdf-output-translations.conf.d/99-disab
```

This assumes that you didn't otherwise configure the ASDF files (if you did, edit them again), and don't somehow override the configuration at runtime with a shell variable (see below) or some other runtime command (e.g. some call to `asdf:initialize-output-translations`).

To disable the compiler output cache in Lisp processes run by your current shell, try (assuming `bash` or `zsh`) (on Unix and cygwin only):

```
export ASDF_OUTPUT_TRANSLATIONS=/:
```

To disable the compiler output cache just in the current Lisp process, use (after loading ASDF but before using it):

```
(asdf:disable-output-translations)
```

12.5 Issues with using and extending ASDF to define systems

12.5.1 "How can I cater for unit-testing in my system?"

ASDF provides a predefined test operation, `test-op`. See Section 6.1.1 [Predefined operations of ASDF], page 14. The test operation, however, is largely left to the system definer to specify. `test-op` has been a topic of considerable discussion on the asdf-devel mailing list, and on the launchpad bug-tracker.

Here are some guidelines:

- For a given system, *foo*, you will want to define a corresponding test system, such as *foo-test*. The reason that you will want this separate system is that ASDF does not out of the box supply components that are conditionally loaded. So if you want to have source files (with the test definitions) that will not be loaded except when testing, they should be put elsewhere.

- The *foo-test* system can be defined in an asd file of its own or together with *foo*. An aesthetic preference against cluttering up the filesystem with extra asd files should be balanced against the question of whether one might want to directly load *foo-test*. Typically one would not want to do this except in early stages of debugging.

- Record that testing is implemented by *foo-test*. For example:

```
(defsystem foo
    :in-order-to ((test-op (test-op foo-test)))
    ....)

(defsystem foo-test
    :depends-on (foo my-test-library ...)
    ....)
```

This procedure will allow you to support users who do not wish to install your test framework.

One oddity of ASDF is that `operate` (see Section 6.1 [Operations], page 14) does not return a value. So in current versions of ASDF there is no reliable programmatic means of determining whether or not a set of tests has passed, or which tests have failed. The user must simply read the console output. This limitation has been the subject of much discussion.

12.5.2 "How can I cater for documentation generation in my system?"

The ASDF developers are currently working to add a `doc-op` to the set of predefined ASDF operations. See Section 6.1.1 [Predefined operations of ASDF], page 14. See also https://bugs.launchpad.net/asdf/+bug/479470.

12.5.3 "How can I maintain non-Lisp (e.g. C) source files?"

See `cffi`'s `cffi-grovel`.

12.5.4 "I want to put my module's files at the top level. How do I do this?"

By default, the files contained in an asdf module go in a subdirectory with the same name as the module. However, this can be overridden by adding a `:pathname ""` argument to the module description. For example, here is how it could be done in the spatial-trees ASDF system definition for ASDF 2:

```
(asdf:defsystem :spatial-trees
  :components
  ((:module base
            :pathname ""
            :components
            ((:file "package")
             (:file "basedefs" :depends-on ("package"))
             (:file "rectangles" :depends-on ("package"))))
   (:module tree-impls
            :depends-on (base)
            :pathname ""
            :components
            ((:file "r-trees")
             (:file "greene-trees" :depends-on ("r-trees"))
             (:file "rstar-trees" :depends-on ("r-trees"))
             (:file "rplus-trees" :depends-on ("r-trees"))
             (:file "x-trees" :depends-on ("r-trees" "rstar-trees"))))
   (:module viz
            :depends-on (base)
            :pathname ""
            :components
            ((:static-file "spatial-tree-viz.lisp")))
   (:module tests
            :depends-on (base)
            :pathname ""
            :components
```

```
                        ((:static-file "spatial-tree-test.lisp")))
            (:static-file "LICENCE")
            (:static-file "TODO")))
```

All of the files in the `tree-impls` module are at the top level, instead of in a '`tree-impls/`' subdirectory.

Note that the argument to `:pathname` can be either a pathname object or a string. A pathname object can be constructed with the '`#p"foo/bar/"`' syntax, but this is discouraged because the results of parsing a namestring are not portable. A pathname can only be portably constructed with such syntax as `#.(make-pathname :directory '(:relative "foo" "bar"))`, and similarly the current directory can only be portably specified as `#.(make-pathname :directory '(:relative))`. However, as of ASDF 2, you can portably use a string to denote a pathname. The string will be parsed as a `/`-separated path from the current directory, such that the empty string `""` denotes the current directory, and `"foo/bar"` (no trailing `/` required in the case of modules) portably denotes the same subdirectory as above. When files are specified, the last `/`-separated component is interpreted either as the name component of a pathname (if the component class specifies a pathname type), or as a name component plus optional dot-separated type component (if the component class doesn't specifies a pathname type).

13 TODO list

Here is an old list of things to do, in addition to the bugs that are now tracked on launchpad:
https://launchpad.net/asdf.

13.1 Outstanding spec questions, things to add

** packaging systems

*** manual page component?

** style guide for .asd files

You should either use keywords or be careful with the package that you evaluate
defsystem forms in. Otherwise (`defsystem partition ...`) being read in the `cl-user`
package will intern a `cl-user:partition` symbol, which will then collide with the
`partition:partition` symbol.

Actually there's a hairier packages problem to think about too. `in-order-to` is not a
keyword: if you read `defsystem` forms in a package that doesn't use ASDF, odd things
might happen.

** extending defsystem with new options

You might not want to write a whole parser, but just to add options to the existing
syntax. Reinstate `parse-option` or something akin.

** document all the error classes

** what to do with compile-file failure

Should check the primary return value from compile-file and see if that gets us any closer
to a sensible error handling strategy

** foreign files

lift unix-dso stuff from db-sockets

** Diagnostics

A "dry run" of an operation can be made with the following form:

```
(traverse (make-instance '<operation-name>)
          (find-system <system-name>)
          'explain)
```

This uses unexported symbols. What would be a nice interface for this functionality?

13.2 Missing bits in implementation

** reuse the same scratch package whenever a system is reloaded from disk

** proclamations probably aren't

** when a system is reloaded with fewer components than it previously had, odd things
happen

We should do something inventive when processing a `defsystem` form, like take the list
of kids and `setf` the slot to `nil`, then transfer children from old to new list as they're found.

** (stuff that might happen later)

*** Propagation of the `:force` option.

"I notice that

`(asdf:compile-system :araneida :force t)`

also forces compilation of every other system the `:araneida` system depends on. This is rarely useful to me; usually, when I want to force recompilation of something more than a single source file, I want to recompile only one system. So it would be more useful to have `make-sub-operation` refuse to propagate `:force t` to other systems, and propagate only something like `:force :recursively`.

Ideally what we actually want is some kind of criterion that says to which systems (and which operations) a `:force` switch will propagate.

The problem is perhaps that "force" is a pretty meaningless concept. How obvious is it that `load :force t` should force *compilation*? But we don't really have the right dependency setup for the user to compile `:force t` and expect it to work (files will not be loaded after compilation, so the compile environment for subsequent files will be emptier than it needs to be)

What does the user actually want to do when he forces? Usually, for me, update for use with a new version of the Lisp compiler. Perhaps for recovery when he suspects that something has gone wrong. Or else when he's changed compilation options or configuration in some way that's not reflected in the dependency graph.

Other possible interface: have a "revert" function akin to `make clean`.

   `(asdf:revert 'asdf:compile-op 'araneida)`

would delete any files produced by `(compile-system :araneida)`. Of course, it wouldn't be able to do much about stuff in the image itself.

How would this work?

`traverse`

There's a difference between a module's dependencies (peers) and its components (children). Perhaps there's a similar difference in operations? For example, `(load "use")` `depends-on (load "macros")` is a peer, whereas `(load "use") depends-on (compile "use")` is more of a "subservient" relationship.

14 Inspiration

14.1 mk-defsystem (defsystem-3.x)

We aim to solve basically the same problems as `mk-defsystem` does. However, our architecture for extensibility better exploits CL language features (and is documented), and we intend to be portable rather than just widely-ported. No slight on the `mk-defsystem` authors and maintainers is intended here; that implementation has the unenviable task of supporting pre-ANSI implementations, which is no longer necessary.

The surface defsystem syntax of asdf is more-or-less compatible with `mk-defsystem`, except that we do not support the `source-foo` and `binary-foo` prefixes for separating source and binary files, and we advise the removal of all options to specify pathnames.

The `mk-defsystem` code for topologically sorting a module's dependency list was very useful.

14.2 defsystem-4 proposal

Marco and Peter's proposal for defsystem 4 served as the driver for many of the features in here. Notable differences are:

- We don't specify output files or output file extensions as part of the system.

  If you want to find out what files an operation would create, ask the operation.

- We don't deal with CL packages

  If you want to compile in a particular package, use an `in-package` form in that file (ilisp / SLIME will like you more if you do this anyway)

- There is no proposal here that `defsystem` does version control.

  A system has a given version which can be used to check dependencies, but that's all.

The defsystem 4 proposal tends to look more at the external features, whereas this one centres on a protocol for system introspection.

14.3 kmp's "The Description of Large Systems", MIT AI Memu 801

Available in updated-for-CL form on the web at http://nhplace.com/kent/Papers/Large-Systems.html

In our implementation we borrow kmp's overall `PROCESS-OPTIONS` and concept to deal with creating component trees from `defsystem` surface syntax. [this is not true right now, though it used to be and probably will be again soon]

Concept Index

Function and Class Index

Variable Index