

CMT

Configuration Management Tool

Version v1r10

Christian Arnault

arnault@lal.in2p3.fr

Contents

1. Copyright.
2. Presentation.
3. The conventions.
4. The architecture of the environment.
 1. Supported platforms
5. Installing a new package.
6. Localizing a package - The **CMPATH** environment variable.
7. Managing site dependent features - The **CMSITE** environment variable.
8. Configuring a package.
9. Selecting a specific configuration.
 1. Describing a configuration.
 2. Defining the user tags.
 3. Activating tags.
10. Working on a package.
 1. Working on a library.
 2. Working on an application
 3. Working on a test or external application
 4. Construction of a global environment
11. Defining a document generator.
 - 1.
 2. How to create and install a new document style
 3. Examples
12. The tools provided by **CMT**
 1. The **requirements** file
 1. The general requirements syntax
 2. The complete requirements syntax
 2. The concepts handled in the requirements file
 1. Meta-information : author, manager
 2. package, version
 3. Constituents : application, library, document
 4. Groups
 5. Languages

6. Symbols : alias, set, set append, set prepend, set remove, macro, macro append, macro prepend, macro remove, macro remove all, path, path append, path prepend, path remove
7. use
8. pattern, apply pattern, ignore pattern
9. branches
10. build strategy, version strategy
11. setup script, cleanup script
12. include path
13. include dirs
14. make fragment
15. public, private
16. tag
3. The general **cmt** user interface
 1. **cmt broadcast [-select=list] [-exclude=list] [-local] [-depth=n] [-all packages] <shell command>**
 2. **cmt build <option>**
 3. **cmt check configuration**
 4. **cmt check files <reference-file> <new-file>**
 5. **cmt checkout ...**
 6. **cmt co ...**
 7. **cmt cleanup [-csh|-sh]**
 8. **cmt config**
 9. **cmt create <package> <version> [<area>]**
 10. **cmt filter <in-file> <out-file>**
 11. **cmt help**
 12. **cmt lock [<package> <version> [<area>]]**
 13. **cmt remove <package> <version> [<area>]**
 14. **cmt remove library links**
 15. **cmt run shell-command**
 16. **cmt setup [-csh|-sh]**
 17. **cmt show <option>**
 18. **cmt system**
 19. **cmt unlock [<package> <version> [<area>]]**
 20. **cmt version**
 21. **cmt cvstags <module>**
 22. **cmt cvsbranches <module>**
 23. **cmt cvssubpackages <module>**
4. The setup and cleanup scripts
5. **cmt build prototype**
13. Using **cv**s together with **CMT**
 1. Importing a package into a **cv**s repository
 2. Checking a package out from a **cv**s repository
 3. Querying **CVS** about some important infos
 4. Working on a package, creating a new release
 5. Getting a particular tagged version out of **cv**s
14. Interfacing an external package with **CMT**
15. Installing **CMT** for the first time

1. Installing CMT on your Unix site
 2. Installing CMT on a Windows or Windows NT site
 16. Differences with previous versions of CMT
 1. Converting a package that was managed with previous versions of CMT (or methods)
 2. Operations in a Windows environment
 17. Appendices
 1. Standard make targets predefined in CMT
 2. Standard macros predefined in CMT
 1. Structural macros
 2. Language related macros
 3. Package customizing macros
 4. Constituent specific customizing macros
 5. Source specific customizing macros
 6. Generated macros
 7. Utility macros
 3. Standard templates for makefile fragments
 4. Makefile generation sequences
 5. The complete requirements syntax
 6. The internal mechanism of cmt cvs operations
-

1 - Copyright.

Copyright (c) 1996 LAL Orsay, UPS-IN2P3-CNRS (France).

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- All advertising materials mentioning features or use of this software must display the following acknowledgement:

This product includes software developed by the
Computer Application Development Group at LAL Orsay
(Laboratoire de l'Accelérateur Linaire - UPS-IN2P3-CNRS).

- Neither the name of the Institute nor of the Laboratory may be used to endorse or promote products derived from this software without specific prior written permission.

This software is provided by the LAL and contributors “as is” and any express or implied warranties, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose are disclaimed. In no event shall the LAL or contributors be liable for any direct, indirect, incidental, special, exemplary, or consequential damages (including, but not limited to, procurement of substitute goods or services; loss of use, data, or profits; or business interruption) however caused and on any theory of liability, whether in contract, strict liability,

or tort (including negligence or otherwise) arising in any way out of the use of this software, even if advised of the possibility of such damage.

2 - Presentation.

This environment, based on some management conventions and comprising several shell-based utilities, is an attempt to formalize software production and especially configuration management around a *package*-oriented principle.

The notion of *packages* represents hereafter a set of software components (that may be applications, libraries, documents, tools etc...) that are to be used for producing a *system* or a *framework*. In such an environment, several persons are assumed to participate in the development and the components themselves are either independent or related to each other.

The environment provides conventions (for *naming* packages, files, directories and for *addressing* them) and tools for *automating* as much as possible the implementation of these conventions. It permits the *description* of the configuration requirements and automatically deduce from the description the effective set of configuration parameters needed to operate the packages (typically for *building* them or *using* them).

CMT lays upon some organisational or managerial principles or mechanisms described below. However, it permits in many respects the users or the managers to *control*, specialize and customize these mechanisms, through parameterization, strategy control and generic specifications.

- Many such packages are produced and maintained.
- The packages may or not be related to each other (defining a *direct acyclic graph* of packages - not just a single tree).
- Each *executable application* (from now on simply named *applications*) either belongs to a particular package and/or defines its own environment and then makes use of some other packages.
- Each package can be uniquely identified within the system or the framework by a *name* which is usually a short *mnemonic* and which may be also used for isolating its name-space (eg. by *prefixing* components of the package by its mnemonic).
- A package installed in this environment may be *exported* to a site where the architecture is reproduced, and as long as the local organisation defined for the package is preserved through the transport, the reconstruction procedure will be preserved. Configuration specifications can be easily provided to cope with machine, site or system specific features.
- Packages are maintained consistently to their declared relationships to each other through a *version* identification model based on :
 - a version is defined with a mnemonic comprising at least two numbers for the *major* id and the *minor* id,
 - versions with different major ids are said to be incompatible,
 - versions with same major ids but different minor ids are said to be backward compatible with respect of the minor id ordering.
- Version control and management schemes (eg. by using **CVS**) are usually consistently operated, applying the conventions on organization and version identification.
- An application that uses one or several packages managed in this environment should not itself be constrained to be managed by **CMT**. The tools should only require a few exported features (such as a few environment variables) for referencing any given package.

- similarly, a package maintained in this environment must be able to use packages that are *not* managed in this environment (which are often called *external* packages).

Following these definitions, the basic configuration management operations involved here (and serviced by the **CMT**’ tools) consist of :

- installing the packages in conventional locations so that they can be referenced by each other, following projects or teams structuring paradigms,
- describing the configuration requirements for each package:
 - dependencies to other packages,
 - symbols to be exported to client packages (environment variables, make macros, etc...)
 - components (also named *constituents*) of the packages (libraries, applications, documents)
 - **Make** macros
 - Strategies that **CMT** should follow at run time, overriding its default ones.
 - Generic behavioural patterns meant to describe generic configuration items.
- deducing the effective configuration parameters from the requirements so as to automatize the building phases and the run-time operations and connections between packages (typically for generating makefiles, generating compiler and linker options, shared libraries paths etc...). This construction mechanism follows customizable strategies (eg. for selecting among possible alternate versions of available packages).

3 - The conventions.

This environment relies on a set of conventions, mainly for organizing the directories where packages are maintained and developed :

- Each package is installed in a standard directory structure defined at least as follows:

```
<some root>/<Package mnemonic>/<version mnemonic>/cmt
```

or / and (*obsolescent convention*)

```
<some root>/<Package mnemonic>/<version mnemonic>/mgr
```

This **cmt** directory holds the main source of information needed by **CMT**: the *requirements* file and all **CMT**-related operations are generally executed from this directory.

This style of organization should be considered as the basic (and unique) criterion for a package to be recognized as a valid *CMT package*.

This structure is a central concept since all relationships between packages relies on the package identification which unambiguously and exclusively consists in the duet [*package-name, package-version*].

- Constructing the internal structure of a package.

Many other parallel branches (similar to **cmt**) such as **src**, **include** or **test** may be freely added to this list according to the specific needs of each package. In particular, a set of such parallel branches are expected to contain *binary* outputs (those that compilers, linkers, archive managers or other kinds of code or pseudo-code generators can produce). Their name always corresponds to the particular *configuration tag* used to produce the output (such as the machine or operating system type). The **CMT** toolkit provides, through the **cmt system** utility, a default value for this token. An environment variable (**CMTCONFIG**) is also assigned to this value (See [this section](#) for a complete description of configuration tags).

Each branch may in addition be freely structured, and there is no constraint to the complexity of this organization.

[IMAGE]

Structuring a package.

- Organizing a software base.

A software base is generally composed of multiple coherent sets of packages, each installed in its specific root directory and forming different *package areas*

There are no constraints on the number of such areas into which **CMT** packages are installed. We'll see [later](#) on how the different areas will be declared and identified by **CMT**.

examples of such organization can be :

[IMAGE]

Structuring a software base.

4 - The architecture of the environment.

This environment is based on the fact that one of its packages (named **CMT**) provides the basic management tools. **CMT**, as a package, has very little specificities and as such itself obeys the general conventions. The major asymetry between **CMT** and all other packages is the fact that once **CMT** is installed it implicitly defines one *default* root area for other packages (through the environment variable **CMTRoot**).

Then packages may be installed either in this default root area or in completely different areas. The only constraint in this case being that their root will have to be specified explicitly.

A typical configuration for this environment consists of selecting a public area (generally available from several machines through an **NFS** or **AFS**-like mechanism), installing the **CMT** basic package, and then installing user packages in this default root or in private ones. One frequent semantic given to this style of configuration is to consider the packages installed in the area hanging below default root as the publicly available versions, whereas packages installed elsewhere are rather meant to be managed in a private context, or in the context of a non public project. However, dependencies between packages will always be possible (as long as the system based protections provide appropriate access rights).

CMT is operated through one main user interface : the **cmt** command, which handles the **CMT** conventions and which provides a set of services for :

- creating a new package, installing it below the default root or in a private area. This operation will create or check the local package directory tree and generate several minimal scripts (see the description of the create command),
- describing or monitoring :
 - the relationships between the package and other packages
 - the configuration features either specified in the current package, or imported from related (used) ones. (symbols, patterns, fragments)
 - the constituents of the package in terms of libraries, executables, or generated documents.
- automatically generating the reconstruction scripts (**makefiles**) from this description.
- recursively acting upon the hierarchy of used packages.

Several other utilities are also provided for some specific activities (such as the automatic production of shared libraries, C prototypes, management of interactions between **CVS** and **CMT** itself, the management of a similar architecture for **Windows** or **OS9**, setting up protections for packages (though locks) etc...).

4.1 - Supported platforms

CMT has been ported and tested on a wide range of machines/operating systems, including :

- DEC-Unix V4.0
- HP-UX-10 (several types of platforms)
- AIX-4
- Solaris
- IRIX
- Several variants of LynxOS
- Linux 2.0
- Windows 95/98/NT/Windows2000 (**nmake** based environment and **MSDev/VisualC++** environment)
- Darwin (Mac OS X)

This in particular means that a package developed on one platform may be re-configured towards any of these platforms without any change to its configuration description (setup scripts, makefiles, ...).

5 - Installing a new package.

We consider here the installation of a user package. Installing **CMT** itself requires special attention and is described in dedicated section of this document.

Therefore, we assume that *some* root directory has been selected by the system manager, and that **CMT** is already installed somewhere. One first has to *setup* **CMT** in order to gain access to the various management utilities, using for example the shell command:

```
csh> source
/lal/CMT/v1r10/mgr/setup.csh
or
ksh> . /lal/CMT/v1r10/mgr/setup.sh
or
dos> call \lal\CMT\v1r10\mgr\setup.bat
```

Obviously, this operation **must** be performed (once) before any other **CMT** action. Therefore it is often recommended to install this setup action straight in the *login* script.

*The setup script used in this example is a constant in the **CMT** environment : every configured package will have one such setup script automatically generated and installed by **CMT**. It is one important entry point to any package (and thus to **CMT** itself). It provides environment variable definitions and recursive invocations of setup scripts for related (used) packages (A corresponding cleanup script is also provided). This script contains a uniform mechanism for interpreting the requirements file so as to dynamically define environment variables, aliases for the package itself and all its used packages. It is constructed once per package installation by the **cmt create** command, or restored by the **cmt config** command (if it has been lost).*

A package is primarily defined by a *name* and a *version* identifier (this duet actually forms the complete *package identifier*). These two attributes will be given as arguments to **cmt create** such as in the following example :

```
csh> cd mydev
csh> cmt create Foo v1
-----
Configuring environment for package Foo version v1.
CMT version v1r10. [1]
Root set to /users/dsksi/arnault/mydev
System is alpha [2]
-----
Installing the package directory [3]
Installing the version directory
Installing the cmt directory
Creating setup scripts.
Creating cleanup scripts.
  1. This shows which actual CMT version you are currently using
  2. This shows the current configuration tag (also available by the cmt system
      command). In this example this is a Compaq alpha machine
  3. This shows the detailed construction of the complete directory structure, starting from
      the top directory which has the name of the package. Since we are creating a completely
      new package, there will be by default only two branches below the version directory :
      cmt and src.
```

The package creation occurred from the current directory, creating from there the complete directory tree for this new package.

In the next example, we install the package in a completely different area, by explicitly specifying the path to it as a third argument to **cmt create** :

```
> cmt create Foo v1 ~/Packages
-----
Configuring environment for package Foo version
v1.
CMT version v1r10.
Root set to /users/dsksi/arnault/Packages.
System is alpha
-----
Installing the package directory
Installing the version directory
Installing the cmt directory
Creating setup scripts.
Creating cleanup scripts.
```

Several file creations occurred at this level :

- a minimal directory tree for the package, including **src** and **cmt** (the other branches will be installed when needed or generated at build time).
- an empty configuration specification file (named **requirements**) installed in the **cmt** branch.
- A minimal **Makefile** (on Unix environments), containing

```
include $(CMTROOT)/src/Makefile.header
include
$(CMTROOT)/src/constituents.make
```

This **Makefile** does not need any modification to build any of the constituents managed by **CMT**. The intermediate makefile fragments will always be re-generated transparently and automatically at build time. However (and thanks to this feature), this file will not be modified anymore by **CMT** itself. Thus you may insert any particular make statement you would feel appropriate, typically when you ask for operations that cannot be taken into account by **CMT**.

- A similar minimal **NMake** file (on Windows environments), containing

```
!include $(CMTROOT)\src\NMakefile.header
!include
$(CMTROOT)\src\constituents.nmake
```

- the setup and cleanup scripts (one flavour for each shell family).

One *may* then setup this new package by running the setup script (which will not have much effect yet since the requirements file is empty) :

```

csh> cd ~/mydev/Foo/v1/cmt
csh> source setup.csh
OR
csh> cd
~/Packages/Foo/v1/cmt
csh> source setup.csh
OR
dos> cd
\Packages\Foo\v1\cmt
dos> call setup.bat

```

The **FOOROOT** and **FOOCONFIG** environment variables are defined automatically by this operation.

It should be noted that running the setup script of a package is not always necessary for building operations. The only situation where running this script *may* become useful, is when an application is to be run, while requiring domain specific environment variables defined in one of the used packages. Besides this particular situation, running the setup scripts may not be needed at all.

Lastly, this newly created package may be removed by the quite similar remove command, using exactly the same arguments as those used for creating the package.

```

csh> cd mydev
csh> cmt remove Foo v1
-----
Removing package Foo version v1.
CMT version v1r10.
Root set to /users/dsksi/arnault/mydev.
System is alpha
-----
Version v1 has been removed from
/users/dsksi/arnault/mydev
Package Foo has no more versions. Thus it has been
removed.

```

OR:

```

csh> cmt remove Foo v1 ~/Packages
-----
Removing package Foo version v1.
CMT version v1r10.
Root set to /users/dsksi/arnault/Packages.
System is alpha
-----
Version v1 has been removed from
/users/dsksi/arnault/Packages
Package Foo has no more versions. Thus it has been removed.

```

So far our package is not very useful since no constituent (application or library) is installed yet. You can jump to the section showing how to work on an [application](#) or on a [library](#) for details on these operations or we can roughly draw the sequence used to specify and build the simplest application we can think of as follows:

```
csh> cd ../src
csh> cat >FooTest.c
#include <stdio.h>
int main ()
{
    printf ("Hello Foo\n");
    return (0);
}
csh> cd ../cmt
csh> vi requirements
...
application FooTest FooTest.c
csh> gmake
csh>
../${CMTCONFIG}/FooTest.exe
Hello Foo
```

This can still be simplified by providing the `-check` option to the application definition:

```
csh> cd ../cmt
csh> vi requirements
...
application FooTest -check
FooTest.c
csh> gmake check
Hello Foo
```

6 - Localizing a package - The CMTPATH configuration parameter.

In the next sections, we'll see that packages *reference* each other by means of *use* relationships. CMT provides a quite flexible mechanism for *localizing* the referenced packages.

A given version of a given package is always referred to by using a *use* statement within its **requirements** file. This statement should specify the package through three *keys* :

- its name (such as **Cm**)
- its version (such as **v7r5**)
- optionally its expected location or prefix (such as **/lal**) (also called the *use path*)

```
use Cm v7r5
or
use Cm v7r5 A
or
use Cm v7r5
/projectB/A
```

Given these keys, the referenced package is looked for according to a prioritized search list which is (in decreasing priority order) :

1. the absolute access path, if the *use path* is absolute,
2. the access paths optionally registered in the configuration parameter - see below - **CMTPATH** (and in decreasing priority, the first element being searched for first),
3. the default root.
4. the path where the current package is installed,

*The configuration parameter **CMTPATH** can be specified either in the environment variable named **CMTPATH** or in **.cmtrc** files, which can themselves be located either in the current directory, in the home directory of the developer or in $\${CMTROOT}/mgr$. In the Windows environment, this configuration parameter may also be installed as a Registry under either the keys:*

- **HKEY_LOCAL_MACHINE/Software/CMT/path**
- **HKEY_CURRENT_USER/Software/CMT/path**

(A graphical tool available in $\%CMTROOT\%VisualC\install.exe$ permits the interactive modification of this list)

If the *path* argument is specified as a relative path (ie. there is no leading *slash* character or it's not a *disk* on windows machines), it will be used as an *offset* to each search case. The search is done starting from the list specified in the **CMTPATH** configuration parameter, then using the default root; and the offset is appended at each searched location.

The **CMTPATH** parameter is thus used as a search list for the packages, and the individual paths are separated in this list by *colons* (*semi-colons* on Windows).

As an example, if we specify the **CMTPATH** parameter as follows :

```
ssh> setenv CMTPATH
/users/dsksi/arnault/dev:/exp/virgo/projects
or (in a requirements file)
path_append CMTPATH "/users/dsksi/arnault/dev"
path_append CMTPATH "/exp/virgo/projects"
or (in a .cmtrc file)
CMTPATH=/users/dsksi/arnault/dev:/exp/virgo/projects
```

Then a *use* statement (defined within a given package) containing :

```
...
use Cm v7r5
use Cmo v1
Cm
...
```

(and assuming that the default root is `/lal`) would look for the package `Cm` from :

1. `/users/dsksi/arnault/dev/Cm/v7r5/cmt`
2. `/exp/virgo/projects/Cm/v7r5/cmt`
3. `/lal/Cm/v7r5/cmt`
4. the same path as the current package

Whereas the package `Cmo` would be searched from :

1. `/users/dsksi/arnault/dev/Cm/Cmo/v1/cmt`
2. `/exp/virgo/projects/Cm/Cmo/v1/cmt`
3. `/lal/Cm/Cmo/v1/cmt`
4. the directory `Cm` within the same path as the current package,

The packages are searched assuming that the directory hierarchy below the access paths always follow the convention :

1. there is a first directory level named according to the package name,
2. then the next directory level is named according to the version tag,
3. then there is a branch named `cmt`,
4. lastly there is a `setup` script within this `cmt` branch.

Thus the list of access paths is searched for until these conditions are properly met.

The actual complete search list can always be visualized by the command:

```
> cmt show path
# Add path /users/dsksi/arnault/dev from CMTPATH
# Add path /exp/virgo/projects from CMTPATH
# Add path /lal from default path
# Add path /tmp/arnault from current package
#
/users/dsksi/arnault/dev:/exp/virgo/projects:/lal:/tmp/arnault
```

7 - Managing site dependent features - The `CMTSITE` environment variable.

Software bases managed by `CMT` are often replicated to multiple geographically distant sites (as opposed to machines connected through AFS-like WAN). In this kind of situation, some of the configuration parameters (generally those used for instance to reference local installations of *external* software) take different values.

The `CMTSITE` environment variable or *registry* in Windows environments, is entirely under the control of the *site* manager and can be set up with a value representing the site (typical values may be `LAL`, `Virgo`, `Atlas`, `LHCb`, `CERN`, etc.).

This variable, when set, corresponds to a *tag* which can be used to select different values for make macros or environment variables.

A typical use for this tag is to build up actual values for the location path of an external software package. Here we take the example of the Anaphe utility:

```
macro AnapheTOP "" \  
    CERN "/afs/cern.ch/sw/lhcxx" \  
    BNL "/afs/rhic/usatlas/offline/external/lhcxx" \  
    LBNL "/auto/atlas/sw/lhcxx"
```

8 - Configuring a package.

The first ingredient of a proper package configuration is the set of configuration parameters which has to be specified in a text file named **requirements** and installed in the **cmt** branch of the package local tree.

An empty version of this file is automatically created the first time the package is installed, and the package manager is expected to augment it with configuration specifications.

Many configuration parameters are supposed to be described into this **requirements** file (one per package) - see the detailed syntax specifications here - namely :

- the package information about its author(s) and manager(s)
- the relationships with other packages
- the package constituents (libraries, applications, documents, etc.)
- the parameterization of the tools used in the build process (eg. make macros)
- the parameterization of the run-time activity (eg. environment variables, search paths, etc.)

Generally, every such appropriate parameter will be deduced on demand from the **requirements** file(s) through the various query functions available from the **cmt** main driver. Therefore there is no systematic package re-configuration per se, besides the very first time a package is newly installed in its location (using the **cmt create** action).

Query actions (generally provided using the **cmt show family** of commands) are embedded in the various productivity tools, such as the setup shell scripts, or makefile fragment generators.

These query actions always interpret the set of **requirements** files obtained from the current package *and* from the packages in the effective *used* chain. Symbols, tags and other definitions are then computed and built up according to inheritance-like mechanisms set up between used packages.

Other configuration parameters are also optionally inserted from the **HOME** and **USER context** requirements files

Most typical examples of these query functions are:

- **cmt setup** builds a shell command line for setting up environment variables
 - **cmt show macros** construct the effective set of inherited make macros
 - **cmt show uses** gives the ordered and flattened set of used packages
 - **cmt show constituents** lists the package's constituents
 - **cmt show path** lists the effective search path for packages.
 - **cmt show strategies** shows the current setup of various functional **CMT** strategies.
-

9 - Selecting a specific configuration.

A configuration describes the conditions in which the package has to be built (ie. compiled and linked) or applications can be run. This configuration can depend on :

- the operating system (such as *Linux*, *Windows*, ...)
- the platform (such as *Intel*, *Compaq*, *Sun*, etc...)
- the choice of the compiler (such as *g++*, *egcs*, *CC*, etc...)
- options used for compiling (such as *optimizer*, *debugger*, etc...) or linking
- the context specifications (selecting a particular version of a firmware, selecting a database server, ...)

Carefully describing this configuration is essential both for maintenance operations (so as to remember the precise conditions in which the package was built) and when the development area is *shared* between machines running different operating systems.

9.1 - Describing a configuration.

CMT relies on several complementary conventions or mechanisms for this description and the associated management.

- The basic binary specification automatically computed by **CMT** in the `${CMTROOT}/mgr/cmt_system.sh` shell script.

This script automatically builds a value characterizing both the machine type and the operating system type (using a mixing of the **uname** standard **UNIX** command with various operating system specific definitions such as the **AFS** based **fs sysname** command)

- The **CMTCONFIG** environment variable, filled in by default from the latter, but which can be manually overridden either from the shell or from the requirements files themselves
- The **CMTSITE** environment variable defines one additional configuration tag, which characterizes the current site.
- The value given by the *uname* standard Unix facility is always a valid configuration tag.
- The concept of user defined *tag set*. Tags are additional qualifiers for the configuration, they are entirely user defined, and have no a-priori semantics.

CMT defines the concept of *current tag set*, as the set of currently active tags. And the current active tag set can always be visualized using the **cmt show tags** command.

9.2 - Defining the user tags.

The user configuration tags can generally be specified through various complementary means:

- CMTSITE and CMTCONFIG can be specified using standard shell commands (setenv, export, set)

```
> export
CMTSITE=CERN
```

- CMTSITE and CMTCONFIG can alternatively be specified using the **set** statement in a requirements file

```
set CMTSITE "CERN"
set CMTCONFIG "${CMTBIN}" sun
"Solaris-CC-dbg"
```

- Additional tags may also be defined as a mixture of other tags, using the **tag** statement (in a requirements file):

```
tag newtag tag1 tag2
tag3
```

which means that:

- **newtag** defines a mixture of tag1 tag2 tag3
- when **newtag** is active, then both tag1, tag2 and tag3 are simultaneously active

9.3 - Activating tags.

By default, only CMTCONFIG, *uname* and CMTSITE are active at any time.

Then it is possible to *activate* alternate tags through the following arguments to *any* cmt command:

- -tag=<primary-tag>

will cleanup the complete current tag set, and provide the new main tag.

- -tag_add=<tag-list>

will add to the current tag set the tags specified in the comma separated list

- -tag_remove=<tag-list>

will remove from the current tag set the tags specified in the comma separated list

Giving these arguments generally make the selected tag set active only during the selected command. However if instead they are given to the **source setup.[c]sh** command, then the new active tag set becomes *persistent* for the current session.

The current active tag set can always be visualized using the **cmt show tags** command.

```
> cmt show tags
Linux
LAL
> cmt -tag_add=tag1,tag2,tag3 show tags
Linux
LAL
tag1
tag2
tag3
> cmt show tags
Linux
LAL
> source setup.sh
-tag_add=tag1,tag2,tag3
> cmt show tags
Linux
LAL
tag1
tag2
tag3
> source setup.sh -tag_remove=tag2,tag3
> cmt show tags
Linux
LAL
tag1
```

Typical usages of those extra tags are:

- when using special compiler options (e.g. optimization, debugging, ...)
- for switching to different compilers (e.g. **gcc** versus the native compiler)
- when one uses a special debugging environment such as **Insure** or **Purify**
- when using special system specific features (such as whether one uses thread-safe algorithms or not)

Due to the tag set concept, it is possible to specify, for instance, that although the current context will still use the default binary tag (ie CMTCONFIG is not changed) a *debug* environment is used.

```
sh> cd .... /Bar/v1/cmt
sh> . setup.sh
-tag_add=debug
```

Then all symbol definitions providing specific values triggered by the **debug** selector will be selected, such as in:

```
macro_append cppflags ""
\
debug " -g "
```

10 - Working on a package.

In this section, we'll see, through a quite simple scenario, the typical operations generally needed for installing, defining and building a package. We are continuing the example of the **Foo** package already used in this document.

10.1 - Working on a library.

Let's assume, as a first example, that the **Foo** package *is* originally composed of one library **libFoo.a** itself made from two sources : **FooA.c** and **FooB.c**. A shared flavour of the library **libFoo.so** or **libFoo.sl** or **libFoo.dll**) is also foreseen.

The minimal set of branches provided by **CMT** (once the **cmt create** operation has been performed) for a package includes **src** for the sources and **cmt** for the *Makefiles* and other scripts.

The various tools **CMT** provide will be fully exploited if one respects the roles these branches have to play. However it is always possible to extend the default understanding **CMT** gets on the package by appropriate modifiers (typically by overriding standard macros).

Assuming the conventional usage is selected, the steps described in this section can be undertaken in order to actually develop a software package.

We first have to create the two source files into the **src** branch (typically using our favourite text editor). Then a description of the expected library (ie. built from these two source files) will be entered into the **requirements** file. The minimal syntax required in our example will be :

```
csh> cd ../cmt
csh> vi requirements (1)
library Foo FooA.c FooB.c
```

1. the requirements file located in the **cmt** branch of the package receives the description of this *library* component. This is done using one **library** statement.

The **cmt create** command did generate a simple *Makefile* (or **NMake** file) which is generally sufficient for all standard operations, since **CMT** continuously and transparently manages the automatic reconstruction of all intermediate makefile fragments. We therefore simply and immediately execute **gmake** as follows:

```

...v1/cmt> [g]make
Rebuilding cmt_path.make [1]
Rebuilding constituents.make [1]
Rebuilding library links [1]
Rebuilding setup.make [1]
alpha.make ok [2]
Library Foo [1][3]
starting Foo [4]
gmake[1]: Entering directory
`/users/dsksi/arnault/mydev/Foo/v1/cmt'
Rebuilding ../alpha/Foo_dependencies.make [1]
gmake[1]: Leaving directory
`/users/dsksi/arnault/mydev/Foo/v1/cmt'
gmake[1]: Entering directory
`/users/dsksi/arnault/mydev/Foo/v1/cmt'
Now rebuilding ../src/FooA.pp
Now rebuilding ../src/FooB.pp
Foo : Protos ok
../alpha/FooA.o
cd ../alpha/; cc -c -I../src/ -I"../src/" -std1 -o FooA.o
../src/FooA.c
../alpha/FooB.o
cd ../alpha/; cc -c -I../src/ -I"../src/" -std1 -o FooB.o
../src/FooB.c
library
cd ../alpha/; ar -clr ../alpha/libFoo.a ../alpha/FooA.o
../alpha/FooB.o
ranlib ../alpha/libFoo.a
cat /dev/null >../alpha/Foo.stamp
cd ../alpha/; /lal/CMT/v1r10/mgr/cmt_make_shlib_common.sh
noextract alpha Foo
-----> Foo : library ok
-----> Foo ok
gmake[1]: Leaving directory
`/users/dsksi/arnault/mydev/Foo/v1/cmt'
all ok.

```

1. The very first time this rebuilding operation occurs, some makefile fragments have automatically been built so as to contain the extended set of Makefile macros deduced from the effective configuration (read from the **requirements** file). These fragments are automatically rebuilt (if needed) each time one of the **requirements** file changes.
2. The directory which is used for the binaries (i.e. the results of compilation or the libraries) has been automatically created by a generic target (**dirs**) which is defined within **[N]Makefile.header**. A new binary directory will be created each time a new value of the **CMTCONFIG** environment variable is defined or a *tag* is provided on the command line to **make**.
3. Each component of the package (be it a particular *library* or a particular *executable*) will have its own *makefile* fragment (named `../${CMTCONFIG}/<name>.[n]mak[e]`). This dedicated *makefile* takes care of filling up the library and creating the shared library (on the systems where this is possible).
4. These dedicated *makefiles* are automatically *executed* from the main one, and the standard make macro **constituents** can be redefined (e.g. in the **requirements** file) so as to customize the building sequence.

or, for nmake:

```
...v1/cmt> nmake /f
nmake
```

This mechanism relies on some conventional *macros* and incremental *targets* used within the specific makefiles. Some are automatically generated, some have to be specified in user packages. It's quite important to understand the list of possible customization macros, since this is the main communication medium between **CMT** and the package manager. See the complete [table](#) of those conventional macro when you want to interact with the standard CMT behaviour.

10.2 - Working on an application

Assume we now want to add a test program to our development. Then we create a **FooTest.c** source, and generate the associated makefile (specifying that it will be an executable instead of a library) :

```
csch> cd ../src csch> emacs FooTest.c ... csch> cd ../cmt csch>
vi requirements ... application FooTest FooTest.c
```

So that we may simply build the complete stuff by running :

```
> [g]make Checking configuration Rebuilding cmt_path.make
Rebuilding constituents.make Rebuilding setup.make
Rebuilding alpha.make alpha.make ok -----> starting FooTest
Application FooTest gmake[1]: Entering directory
`/users/dsksi/arnault/mydev/Foo/v1/cmt Rebuilding
FooTest_dependencies.make gmake[1]: Leaving directory
`/users/dsksi/arnault/mydev/Foo/v1/cmt' gmake[1]: Entering
directory `/users/dsksi/arnault/mydev/Foo/v1/cmt' Now
rebuilding ../src/FooTest.pp FooTest : Protos ok cd
../alpha/; cc -c -I../src/ -I../src/ -std1 ../src/FooTest.c
cd ../alpha/; cc -o FooTest.exe.new ../alpha/FooTest.o ; \
mv -f FooTest.exe.new FooTest.exe -----> FooTest ok
gmake[1]: Leaving directory
`/users/dsksi/arnault/mydev/Foo/v1/cmt' -----> all ok.
```

Which shows that a program **FooTest.exe** has been built from our sources. Assuming now that this program needs to access the **Foo** library, we'll just add the following definition in the [requirements](#) file :

```
... macro Foo_linkopts " -L$(FOOROOT)/$(Foo_tag) -lFoo " \
WIN32 " $(FOOROOT)/$(Foo_tag)/Foo.lib " ...
```

The **Foo_linkopts** conventional macro will be automatically inserted within the **use_linkopts** macro.

Like all other make macros used to build a component, the **Foo_linkopts** will be specified within the **requirements** which gives several benefits:

- variants of the macro definition can be provided
- monitoring features of **CMT** such as the **cmt show macro Foo_linkopts** command can be used later on
- macros defined this way may be later on inherited by client packages which will *use* our package.

10.3 - Working on a test or external application

It is also possible to work on a *test* or *external* application, ie. when one does not wish to configure the development for this application using **CMT**. Even in this case, it is possible to benefit from the packages configured using **CMT** by partially using **CMT**, just for *used* relationships.

Here, no special convention is assumed on the location of the sources, the binaries, the management scripts, etc... However, it is possible to describe in a **requirements** file the *use* relationships, as well as the **make** macro definitions, quite similarly to the package entirely configured using **CMT**.

Most of the options provided by the **cmt** user interface are still available in these conditions.

10.4 - Construction of a global environment

A software base generally consists in many *packages*, some of them providing *libraries* or *documents*, others providing *applications*, some providing both, some providing just *glues* towards external software products.

On another view, this software base may a mix of packages shared between several projects and sets of packages specific to various projects. One may have several software bases as well (combined using the **CMTPATH** environment variable).

In such contexts, it is often desirable that a given project defines its own selection of all existing packages. This can easily be done with **CMT** by defining a *project* package, containing only **use** statements towards the appropriate selection of packages for this particular project.

Let's consider as an example the project named **MyProject**. We may create the package named **MyProject** similarly to any other package :

```
csH> cd ..... csH> cmt create MyProject
v1
```

Then the **requirements** file of this new package will simply contain a set of **use** statements, defining the *official* set of validated versions of the packages required for the project. This mechanism also represents the notion of *global release* traditionally addressed in configuration management environments

```
package MyProject use Cm v7r6 use Db v4r3 use El v4r2 use Su
v5 use DbUI v1r2 Db use ElUI v1r1 El use VSUUI v3 Su/VSU use
VMM v1 use VPC v3
```

Then any user wanting to access the so-called *official* release of the package set appropriate to the project **MyProject** will simply do (typically within its login shell script) :

```
# a login script ... source
/MyProjectDevArea/MyProject/v1/cmt/setup.csh
```

Later on, future evolutions of the **MyProject** package will reflect progressive integration steps, which *validate* the evolutions of each referenced package.

11 - Defining a document generator

In a Unix environment, documents are built using **make** (well generally its *gnu* flavour) or **nmake** in Windows environments. The basic mechanism provided in **CMT** relies on *make fragment patterns* containing instructions on how to rebuild document pieces. Many such generators are provided by **CMT** itself so as to take care of the most usual cases (e.g. compilations, link operations, archive manipulations, etc...). In addition to those, any package has to possibility to provide a new generator for its own purpose, i.e. either for providing rules for a special kind of document, or even to override the default ones provided by **CMT**. This mechanism is very similar to the definition or re-definition of *macros* or environment variables in that every new generator has to be first declared in a **requirements** file belonging to a package (**CMT** actually declares all its default generators within its **requirements** file), allowing all its client packages to transparently acquire the capacity to generate documents of that sort.

CMT manages two categories of constituents:

1. *Applications* and *Libraries* are handled using pre-defined make fragments (mainly related with languages) and behaviour.
2. *Documents* offer a quite general framework for introducing completely new behaviours through user-defined make fragments. This includes actually generating documents, but also simply performing an operation (in which case sometimes no real *document* is produced).

In this section we only discuss the latter category and the following paragraphs explain the framework used for defining new document types.

The main concept of this framework is that each document to be generated or manipulated must be associated with a "document-type" (also sometimes named "document-style"), which corresponds to a dedicated make fragment of that name. Then, when specified in a **document** statement, this make fragment will be *instanciated* once or several times (typically once per source file) to construct a complete and functional make fragment, containing one main target. Both the resulting make fragment and the make target will have the name of the constituent.

11.1 -

11.2 - How to create and install a new document style

This section presents the general framework for designing a document generator.

1. Select a name for the document style. It should not clash with existing ones (use the `cmt show fragments` for a complete list of document types currently defined).
2. A fragment exactly named after the document style name must be installed into a subdirectory named `fragments` below the `cmt` branch of a given package (which becomes the *provider* package).
3. Optionally, two other fragments may be installed into the same subdirectory, one of them will be the *header* of the generated complete fragment, the other will be its *trailer*
4. Those fragments *must* be declared in the requirements file of the provider package as follows:

```
make_fragment <fragment-name> [ options...  
]
```

where options may be :

<code>-suffix=<suffix></code>	provide the suffix of the output files (without the dot)
<code>-header=<header></code>	provide another make fragment meant to be prepended to the constituent's make fragment.
<code>-trailer=<trailer></code>	provide another make fragment meant to be appended to the constituent's make fragment.
<code>-dependencies</code>	install the automatic generation of dependencies into the constituent's make fragment

Once a fragment is installed and declared, it may be used by any *client* package (ie a package *using* the provider), and queried upon using the command

```
> cmt show fragment <fragment
name>
```

which will show where this fragment is defined (ie. in which of the used packages).

The `cmt show fragments` commands lists all declared fragments.

If a package re-defines an already declared make fragment, ie it provides a new copy of the fragment (possibly with new copies of the header and the trailer), and declares it inside its requirements file, then this package becomes the new provider for the document style.

For building a fragment, one may use pre-defined generic "templates" (which will be substituted when a fragment is copied into the final constituent's makefile).

CONSTITUENT	the constituent name
CONSTITUENTSUFFIX	the optional constituent's output suffix
FULLNAME	the full source path name (including directory and suffix)
FILENAME	the complete source file name (only including the suffix)
NAME	the short source file name (without directory and suffix)
FILEPATH	the source directory
SUFFIX	the suffix provided in the <code>-suffix</code> option
OBJS	(only available in headers) the list of outputs, formed by a set of expressions : <div style="border: 1px solid black; padding: 5px; margin: 5px auto; width: fit-content;"> $\\$(\\${CONSTITUENT}_output)\\${NAME}\\${SUFFIX}$ </div>

Templates must be enclosed between `{` and `}` or between `(` and `)` and will be substituted at the generation time. Thus, if a fragment contains the following text :

```
$( ${CONSTITUENT}_output ) ${NAME} ${SUFFIX}
```

then, the expanded constituent's makefile will contain (referring to the "tex" example)

```
$(MyDoc_output) doc1.ps
```

Which shows that make macros may be dynamically generated.

[IMAGE]

The architecture of document generation.

11.3 - Examples

1. rootcint

It generates C++ hubs for the Cint interpreter in Root.

```
===== rootcint
===== $(src)${NAME}.cc
:: ${FULLNAME} ${rootcint} -f $(src)${NAME}.cc -c
${FULLNAME}
=====
```

2. agetocxx and agetocxx_header.

It generates C++ source files (xxx.g files) from Atlas' AGE description files.

```

===== agetocxx
=====
output=$((${CONSTITUENT}_output) ${output})${NAME}.cxx :
${(${NAME}_cxx_dependencies) (echo '#line 1 "${FULLNAME}";
cat ${FULLNAME}) > /tmp/${NAME}.gh.c gcc -E -I${output}
${use_includes} -D_GNU_SOURCE \ cd ${output}; $(agetocxx)
-o ${NAME} -ohd ${FILEPATH} \ -ohp ${FILEPATH}
/tmp/${NAME}.gh rm -f /tmp/${NAME}.gh /tmp/${NAME}.gh.c cd
$(bin); $(cppcomp) ${use_cppflags}
${(${CONSTITUENT}_cppflags) \ ${(${NAME}_cppflags)
${ADDINCLUDE} ${output})${NAME}.cxx cd $(bin); $(ar)
${(${CONSTITUENT}lib) ${NAME}.o; /bin/rm -f ${NAME}.o
=====

```

```

===== agetocxx_header
===== ${CONSTITUENT}lib =
$(bin)lib${CONSTITUENT}.a ${CONSTITUENT}stamp =
$(bin)${CONSTITUENT}.stamp ${CONSTITUENT}shstamp =
$(bin)${CONSTITUENT}.shstamp ${CONSTITUENT} :: dirs
${CONSTITUENT}LIB @/bin/echo ${CONSTITUENT} ok
${CONSTITUENT}LIB :: ${(${CONSTITUENT}lib)
${(${CONSTITUENT}shstamp) @/bin/echo ${CONSTITUENT} :
library ok ${(${CONSTITUENT}lib) ${(${CONSTITUENT}stamp) ::
${OBJS} $(ranlib) ${(${CONSTITUENT}lib) cat /dev/null
>${(${CONSTITUENT}stamp) ${(${CONSTITUENT}shstamp) ::
${(${CONSTITUENT}stamp) cd $(bin); $(make_shlib) $(tag)
${CONSTITUENT} \ ${(${CONSTITUENT}shlibflags); \ cat
/dev/null >${(${CONSTITUENT}shstamp)
=====

```

It must be declared as follows :

```

make_fragment agetocxx -suffix=cxx -dependencies
-header=agetocxx_header

```

12 - The tools provided by CMT

The set of conventions and tools provided by CMT is mainly composed of :

- **the syntax of the requirements file,**
- **and the general cmt user interface, available in the mgr branch of the CMT package.**

The *setup* script found in the CMT installation directory

actually adds its location to the definition of the standard UNIX PATH environment variable in order to give direct access to the main cmt user interface.

The sections below will detail the complete syntax of the requirements file since it is the basis of most information required to run the tools as well as the main commands available through the cmt user interface.

12.1 - The requirements file

12.1.1 - The general requirements syntax

- A requirements file is made of *statements*, each describing one named configuration parameter.

Statements generally occupy one single line, but may be split into several lines using the reverse-slash character (in this case the reverse-slash character *must* be the last character on the line or must be only followed by space characters).

Each statement is composed of words separated with spaces or tabulations.

The first word of a statement is the name of the configuration parameter.

The rest of the statement provides the value assigned to the configuration parameter.

- Words composing a statement are separated with space or tab characters. They may also be enclosed in quotes when they have to include

space or tab characters. Single or double quotes may be freely used, as long as the same type of quote is used on both sides of the word.

Special characters (tabs, carriage-return and line-feed) may be inserted into the statements using an XML-based convention:

tabulation	<cmt:tab/>
carriage-return	<cmt:cr/>
line-feed	<cmt:lf/>

- **Comments : they start with the # character and extend up to the end of the current line.**
-

12.1.2 - The complete requirements syntax

12.2 - The concepts handled in the requirements file

12.2.1 - Meta-information : author, manager

The author and manager names

12.2.2 - package, version

The package name and version. These statements are purely informational.

12.2.3 - Constituents : application, library, document

Describe the composition of a constituent.

Application and library correspond to the standard meaning of an application (an executable) and a library, while document provides for a quite generic and open mechanism for describing any type of document that can be generated from sources.

Applications and libraries are assigned a name (which will correspond to a generated make fragment, and a dedicated make target).

A document is first associated with a document type (which must correspond to a previously declared make fragment). The document name is then used to name a dedicated make fragment and a make target.

Various options can be used when declaring a constituent:

<i>option</i>	<i>validity</i>	<i>usage</i>
-windows	applications	When used in a Windows environment, generates a GUI-based application (rather than a console application)
-no_share	libraries	do not generate the shared library
-no_static	libraries	do not generate the static library (<i>not yet implemented</i>)
-prototypes	applications, libraries	do generate the prototype header files
-no_prototypes	applications, libraries	do not generate the prototype header files
-check	applications	generate a check target meant to execute the rebuilt application
-group=group-name	any	install the constituent within this group target
-suffix=suffix	applications, libraries	provide a suffix to names of all object files generated for this constituent <u>(1)</u>
-import=package	applications, libraries	explicitly import for this constituent the standard macros from a package that has the <code>-no_auto_imports</code> option set
variable-name=variable-value	any	define a variable and its value to be given to the make fragment <u>(2)</u>

1.

When several constituents need to share source files, (a typical example is for building different libraries from the same sources but with different compiler options), it is possible to specify an optional output suffix with the `-suffix=<suffix>` option. With this option, every object file name will be automatically suffixed by the character string "`<suffix>`", avoiding name conflicts between the different targets, as in the following example:

```
library AXt -suffix=Xt *.cxx library AXaw -suffix=Xaw
*.cxx
```

2.

It's possible to specify in the list of parameters one or more pairs of variable-name=variable-value (without any space characters around the "=" character), such as in the next example:

```
make_fragment doc_to_html (1) document doc_to_html Foo
output=FooA.html FooA.doc (2) (3)
```

1. This makefile fragment is meant to contain some text conversion actions and defines a document type named `doc_to_html`.
2. This constituent exploits the document type `doc_to_html` to convert the source `FooA.doc` into an html file.
3. The user defined template variable named `output` is specified and assigned the value `FooA.html`. If the fragment `doc_to_html` contains the string `${output}`, then it will be substituted to this value.

12.2.4 - Groups

Groups permit the organization of the constituents that must be consistently built at the same development phases or with similar constraints.

Each group is associated with a make target (of the same name) which, when used in the make command, selectively rebuilds all constituents of this

group.

The default group (into which all constituents are installed by default) is named *all*, therefore, running `make` without argument, activates the default target (ie. *all*).

As a typical usage of this mechanism, one may exemplify the case in which one or several constituents are making use of one special facility (such as a database service, real-time features, graphical libraries) and therefore might require a controlled re-build. This is especially useful for having these constituents only rebuilt on demand rather than rebuilt automatically when the default `make` command is run.

One could, for instance specify within the requirements file :

```
# Constituents belonging to the default all group ...  
constituents without group specification ... # Constituents  
belonging to specific groups library Foo-objy -group=objy  
<sources making use of Objectivity> application FooGUI  
-group=graphics <sources making use of Qt> application  
BarGUI -group=graphics <sources making use of Qt>
```

(Beware of the position of the `-group` option which must be located after the constituent name. Any other position will be misunderstood by CMT)

Then, running `gmake all` would only rebuild the un-grouped constituents, whereas running

```
> gmake objy > gmake  
graphics
```

in the context of the Foo package would rebuild *objy* related or *graphics* related constituents.

12.2.5 - Languages

Some computer languages are known by default by CMT (C, C++, Fortran77, Java, lex, yacc). However it is possible to extend this knowledge to any other language.

We consider here languages that are able to produce object files from sources.

Let's take an example. We would like to install support for Fortran90. We first have to *declare* this new language support to CMT within the requirements file of one of our packages (Notice that it's not at all required to modify CMT itself since all clients of the selected package will inherit the knowledge of this language).

The language support is simply named **fortran90** and is declared by the following statement:

```
language fortran90 \ -suffix=f90 -suffix=F90 \ [1]
-linker=$(f90link) \ [2] -preprocessor_command=$(ppcmd)
```

1. The recognized suffixes for source files will be f90 and F90
2. The linker command used to build a Fortran90 application is described inside the macro named f90link (which must be defined in this requirements file but which can also be overridden by clients)

The language support being named **fortran90**, two associated make fragments are expected, one under the name **fortran90** (for building application modules), the other with the name **fortran90_library** (for modules meant to be archived), both without extension.

These two fragments should be installed in the fragments sub-directory of the cmt branch of our

package.

Due to the similarity of the example to `fortran77`, we may easily provide the expected fragments simply by copying the `f77` fragments found in CMT (thus the fragments ``${CMROOT}/fragments/fortran` and ``${CMROOT}/fragments/fortran_library`

These fragments make use of the `fcomp` macro, which holds the `fortran77` compiler command (through the `for` macro).

```
macro for "f77" \ ... macro fcomp "$(for) -c $(fincludes)  
$(fflags) $(pp_fflags)"
```

We therefore simply replace these macros by new macros named `f90comp` and `f90`, defined as follows:

```
macro f90 "f90" ... macro f90comp "$(f90) -c $(fincludes)  
$(fflags) $(pp_fflags)"
```

Some languages (this has been seen for example in the IDL generators in Corba environments) do provide several object files from one unique source file. It is possible to specify this feature through the (repetitive) `-extra_output_suffix` option like in:

```
language idl -suffix=idl -fragment=idl  
-extra_output_suffix=_skel
```

where, in this case, two object files are produced for each IDL source file, one named `<name>.o` the other named `<name>_skel.o`.

12.2.6 - Symbols : `alias`, `set`, `set_append`, `set_prepend`, `set_remove`, `macro`, `macro_append`, `macro_prepend`, `macro_remove`, `macro_remove_all`, `path`, `path_append`, `path_prepend`, `path_remove`

The alias keyword is translated into a shell alias definition,

The set keyword is translated into an environment variable definition.

The macro keyword is translated into a make's macro definition.

The path keyword is translated into a *path*-like environment variable, which is supposed to be composed of search paths separated with colon characters ':' (on Unix) or semi-colon characters ';' (on Windows). However, it is highly recommended to construct such a variable by iteratively concatenating individual items one by one using *path_append* or *path_prepend*

Variants of these keywords are also provided for modifying already defined symbols. This generally happens when a package needs to modify an inherited symbol (ie. which has been already defined by a used package). Through the keywords *set_append*, *set_prepend*, *set_remove*, *macro_append*, *macro_prepend*, *macro_remove*, *macro_remove_all*, *path_append*, *path_prepend*, *path_remove* one can append or prepend a text to the existing symbol value or remove a part from it. The *path_remove* keyword removes all individual search paths that *contain* the specified sub-string.

The translations occur while running either the setup scripts (for alias, set or path) or the make command (for macro).

All these definitions follow the same pattern:

```
<symbol-type> <symbol-name> <default-value> [ <tag> <value>  
... ]
```

The symbol-name identifies the symbol for modification operations. The default-value is optionally followed by a set of tag/value pairs, each representing an alternate value for this symbol. Be aware that there is only one name space for all kinds of symbols. Therefore, if a symbol was originally defined using a macro statement, using set_append to modify it will produce an undefined result.

The tag is used to select one alternate value to replace the default value, when one of the following condition is met:

- **It matches the value of the CMTSITE environment variable (or registry)**
- **It matches the value provided by the uname Unix command (when available)**
- **It matches the value of the CMTCONFIG environment variable (or registry)**
- **It matches the value specified in the -tag=tag argument to the cmt command.**
- **It matches one user defined tag (see the tag keyword) which itself is associated with a matching tag (Note that this is a recursive definition).**

Examples of such definition are :

```
package CMT set CMTCC "cc" \ HP-UX "cc -Aa +z
-D_HPUX_SOURCE" public macro cflags "-g" \ HP-UX "-g -Aa +z
-D_HPUX_SOURCE" \ hp700_ux101 "-g -fpic -ansi" \ alpha "-g
-std1" \ alphas "-g -std1 -DTHREADS" \ insure "-g -Zuse
-std1" \ AIX "-g -D_ALL_SOURCE -D_BSD" macro cppflags "-g"
\ HP-UX "-g -Aa +z" \ hp700_ux101 "-g -fpic" macro fflags
"-g" macro src "../src/" macro inc "../src/" macro mgr
"../cmt/" macro SHELL "/bin/sh"
```

12.2.7 - use

Describe the relationships with other packages; the generic syntax is :

```
use <package> [ <version> [ <root> ]
]
```

Omitting the version specification means that the most recent version (ie. the one with highest ids) that can be found from the search path list will be automatically selected.

The root specification can be relative (ie. on Unix it does not contain a leading '/' character). In this case, this prefix is systematically considered when the package is looked for in the search path list. But it can also be absolute (ie. with a leading '/' character on Unix), in which case this path takes precedence over the standard search path list (see CMTPATH).

Examples of such relationships are :

```
# Packages installed in the default root : use OnX v5r2 use
CSet v2r3 use Gb v2r1 # A package installed in a root one
step below the root : use CS v3r1 virgo # Back to the
default root : use Cm v7r3 # Get the most recent version of
CERNLIB use CERNLIB
```

By default, a set of standard macros, which are expected to be specified by used packages, is

automatically imported from them (see the detailed list of these macros). This automatic feature can be discarded using the `-no_auto_imports` option to the use statement, or re-activated using the `-auto_imports`. When it is discarded, the macros will not be transparently inherited, but rather, each individual constituent willing to make use of them will have to explicitly import them using the `-import=<package>` option.

When a use statement is in a private section, the corresponding used package will only be reached if when CMT operations occur in the context of the holder package. Otherwise (ie if the operation occurs in some upper level client package, then this *privately* used package will be entirely hidden. (*This behaviour follows a very similar pattern to the private or public inheritance of C++*). Suppose we have the following organization:

```

----- package A use B v1 use D v1
----- package B private use C
v1 use D v1 -----

```

- all operations done in the context of package B will *see* both packages C and D
- all operations done in the context of package A will *see* both packages B and D, but not package C

12.2.8 - `pattern`, `apply_pattern`, `ignore_pattern`

Often, similar configuration items are needed over a set of packages (sometimes over all packages of a project). This reflects either similarities between packages or generic conventions established by a project or a team.

Typical examples are the definition of the search path for shared libraries (through the `LD_LIBRARY_PATH` environment variable), the systematic production of test applications, etc.

The concept of pattern proposed here implements this genericity. Patterns may be either *global*, in which case they will be systematically applied onto every package, or *local* (the default) in which case they will be applied on demand only by each package.

The general principle of a pattern is to associate a templated (set of) cmt statement(s) with the pattern name. Then every time the pattern is applied, its associated statements are applied as if they were directly specified in the requirements file, replacing the template with its current value. If several statements are to be associated with a given pattern, they will be separated with the " ; " separator pattern (beware of really enclosing the ";" between two space characters).

Pattern templates are names enclosed between the '<' and '>' characters. A set of predefined templates are automatically provided by CMT:

package	the name of the current package
PACKAGE	the name of the current package in upper case
version	the version tag of the current package
path	the access path of the current package

Then, in addition, user defined templates can be installed within the pattern definitions. Their actual value will be provided as arguments to the `apply_pattern` statement.

User defined templates that have not been assigned a value when the pattern is applied are simply ignored.

Some examples:

1. Changing the standard include search path.

The standard include path is set by default to `${<package>_root}/src`. However, often projects need to override this default convention, and typical example is to set it to a branch named with the package name. This convention is easily applied by defining a pattern which will apply the `include_path` statement as follows:

```
pattern -global include_path include_path
${<package>_root}/<package>/
```

For instance, a package named PackA will expand this pattern as follows:

```
include_path
${PackA_root}/PackA/
```

2. Providing a value to the LD_LIBRARY_PATH environment variable

On some operating systems (eg. Linux), shared library paths must be explicated, through an environment variable. The following pattern can automate this operation:

```
pattern ld_library_path \ path_remove LD_LIBRARY_PATH
"/<package>/" ; \ path_append LD_LIBRARY_PATH
${<PACKAGE>ROOT}/${CMTCONFIG}
```

In this example, the pattern was not set global, so that only packages actually providing shared

libraries would be concerned. These packages will simply have to apply the pattern as follows:

```
apply_pattern
ld_library_path
```

The schema installed by this pattern provides first a cleanup of the `LD_LIBRARY_PATH` environment variable and then the new assignment. This choice is useful in this case to avoid conflicting definitions from two different versions of the same package.

3. Installing a systematic test application in all packages

Quality assurance requirements might specify that every package should provide a test program. One way to enforce this is to build a global pattern declaring this application. Then every make command would naturally ensure its actual presence.

```
pattern quality_test application <package>test
<package>test.cxx <other_sources>
```

In this example, an additional pattern (`<other_sources>`) permits the package to specify extra source files to the test application (the pattern assumes at least one source file `<package>test.cxx`).

12.2.9 - branches

Describe the specific directory branches to be added while configuring the package.

```
branches <branch-name>
...
```

These branches will be created (if needed) at the same level as the cmt branch. Typical examples of such required branches may be include, test or data.

12.2.10 - build_strategy, version_strategy

Users can control the behaviour of CMT through a set of strategy specifications. The current implementation only provides such control over two mechanisms :

- **the way version tags are interpreted and compared to each other.**

The following keywords are available:

best_fit	This is the default behaviour. Version tags truly consider major ids, minor ids and patch ids with their complete backward compatibility semantics
best_fit_no_check	Same as best_fit except that different major ids are not seen as incompatible. The greatest id (for major, minor and patch ids) is always selected
first_choice	The first version tag specified in the use chain is selected
last_choice	The last version tag specified in the use chain is selected
keep_all	Internal use only : all referenced versions are kept

- **the way makefile fragments for applications and libraries are generated.**

Currently this only concerns the automatic generation of prototype header files for C source files. Thus only one keyword is possible : prototypes (and its opposite no_prototypes), the default CMT behaviour being to generate prototype headers.

12.2.11 - `setup_script`, `cleanup_script`

Specify user defined configuration scripts, which will be activated together with the execution of the main setup and cleanup scripts.

The script names may be specified without any access path specification, in this case, they are looked for in the `cmt` or `mgr` branch of the package itself. They may also be specified without any `.csh` or `.sh` suffix, the appropriate suffix will be appended accordingly when needed. Therefore, when such a user configuration script is specified, CMT expects that the corresponding shell scripts actually exist in the appropriate directory (the `cmt` branch by default) and is provided in whatever format is appropriate (thus suffixed by `.csh` and/or `.sh`).

12.2.12 - `include_path`

Override the specification for the default include search path, which is internally set to `${<package>_root}/src`.

Specifying the value *none* (a reserved CMT keyword) means that no default include search path is expected from CMT, and thus no `-I` compiler option will be generated by default (generally this means that user include search paths should be specified via *include_dirs* instead).

12.2.13 - `include_dirs`

Add specifications for non-standard include access paths.

12.2.14 - make_fragment

This statement specifies a specialized makefile fragment, used as a building brick to construct the final makefile fragment dedicated to build the constituents.

There are basically three categories of such fragments :

1. some are provided by CMT itself (they correspond to its internal behaviour)
2. others handle the language support
3. and the last serve as specialized document generators.

The fragments defined in CMT can be:

- those used to construct the application or library constituents. Their semantic is standardized (they are all associated with a language statement in the CMT requirements file).

*c c_library cpp cpp_library lex lex_library
fortran fortran_library yacc yacc_library jar
jar_header java java_copy java_header
check_java cleanup_java*

- those used internally by CMT as primary building blocks for assembling the makefile. (Generally developers should not see them).

*cleanup_objects application make_setup_nmake
constituent application_header
constituents_header buildproto protos_header
os9_header dependencies check_application
dependencies_and_triggers*

*check_application_header document_header
library_cleanup library_header cleanup_application
library_no_share cleanup_header make_header make_setup
cleanup_library make_setup_header*

- some document generators which *may* be used if needed, but are not mandatory:

*installer installer_header readme
readme_header readme_trailer readme_use dvi
tex generator generator_header*

- those used to generate configuration files for MSVisualC++:

*dsp_windows_header dsw_all_project
dsw_all_project_dependency
dsw_all_project_header dsw_all_project_trailer
dsw_header dsw_project dsw_trailer dsp_all
dsp_application_header dsp_contents
dsp_library_header dsp_shared_library_header
dsp_trailer*

Language fragments should provide two forms, one for the applications (in which case they are named exactly after the language name eg c, cpp, fortran) and the other for the libraries (in which case they have the `_library` suffix (eg. `c_library`, `cpp_library`, `fortran_library`). A set of language definitions (C, C++, Fortran, Java, Lex, Yacc) is provided by CMT itself but it is expected that projects add new languages according to their needs. Even if the make fragment meant to be the implementation of a language support is declared, the language support itself must be declared too, using the language

statement

All make fragments are provided as (suffixless) files which must be located in the fragments sub-directory inside the cmt/mgr branch of one package. They must also be declared in the requirements file (through the `make_fragment` statement) so as to be visible.

A package declaring, and implementing a make fragment may override a fragment of the same name when it is already declared by a used package. This implies in particular that any package may freely override any make fragment provided by CMT itself (although in this case a deep understanding of what the original fragment does is recommended).

Makefile fragments may take any form convenient to the document style, and some special pre-built templates (see the [appendix](#)) can be used in their body to represent running values, meant to be properly expanded at actual generation time :

CONSTITUENT	the constituent name
FULLNAME	the full source path
FILENAME	the source file name without its path
NAME	the source file name without its path and suffix
FILESUFFIX	the dotted file suffix
FILEPATH	the output path
SUFFIX	the default suffix for output files

12.2.15 - public, private

Introduce a section for *public* or *private* symbols (meant to be implemented as environment variables or aliases in a Unix environment or as *logical names*

or *symbols* in a VMS one). *Macros* to be used within makefiles can also be defined at this level. Public symbols are meant to be exported to any external user of the package whereas private ones are only defined for the package *developer*. Currently the selection between these two categories is done when the setup script is executed : if it is executed while actually being in the `cmt` branch of the package, the *developer* category is assumed. If the script is executed from another directory the *user* category is assumed.

12.2.16 - tag

Provide tag definitions.

A tag is a token which can be used to select particular values of symbols. Generally a tag need not being explicitly declared, since the reference to it will declare the tag automatically. However, tags may be used to *name* a particular association of several other tags. In this case, this association must be declared within a requirements file as follows :

```
tag <association-tag-name> <tag1> <tag2> ... eg: tag
Linux-gcc Linux gcc
```

This definition implies that when *Linux-gcc* is true, then both *Linux* and *gcc* are true.

This can be exploited to trigger via only one tag, the activation of several individual tags, each signing a particular condition (in our example the *debug* condition and the *Linux* environment).

However, it is always possible to dynamically associate several tags by using the *tag-list*-style of

arguments to the **-tag=<tag-list>** options to the **cmt** command driver (such as in *cmt setup -tag=Linux,debug*)

Tags or associations of tags are propagated using the **-tag=<tag-list>** options to the **cmt** command driver, but the **Make** command can also accept them through the conventional macros $$(tag)$ and $$(extra_tags)$.

However, the $$(tag)$ macro itself can only accept one value (instead of a list), and therefore in order to give a list of additional tags, one should use the $$(extra_tags)$ (such as in *gmake tag=Linux extra_tags=debug*)

Finally, running the setup script (through the *source setup.[c]sh* or *call setup.bat* command) can also receive tag specifications using the **-tag=tag-list** options.

12.3 - The general cmt user interface

This utility (a shell script combined with a C application) provides a centralised access to various commands to the CMT system. The first way to use **cmt** is to run it without argument, this will print a minimal help text showing the basic commands and their syntax :



The following sections present the detail of each available command.

12.3.1 - `cmt broadcast [-select=list] [-exclude=list] [-local] [-global] [-begin=pattern] [-depth=<n>] [-all_packages] <shell command>`

This command tries to repeatedly execute a shell command in the context of each of the used package of the current package. The used packages are listed using the `cmt show uses` command, which also indicates in which order the broadcast is performed. When the `all_packages` option, the set of packages reached by the broadcast is rather the same as the one shown by the `cmt show packages` command, ie all CMT packages and versions available through the current `CMPATH` list.

Typical uses of this *broadcast* operation could be:

```
csh> cmt broadcast cmt config csh> cmt broadcast - gmake
csh> cmt broadcast '(cd ../; cvs -n update)'
```

The loop over used packages will stop at the first error occurrence in the application of the command, except if the command was preceded by a '-' (minus) sign (similarly to the make convention).

It is possible to specify a list of selection or exclusion criteria set onto the package path, using the following options, right after the broadcast keyword. These selection criteria may be combined (eg one may combine the *begin* and *select* modifiers)

```
sh> cmt broadcast -begin=Cm gmake (1) sh> cmt broadcast
-select=Cm gmake (2) sh> cmt broadcast -select='/Cm/
/CSet/' gmake (3) sh> cmt broadcast -select=Cm -exclude=Cmo
gmake (4) sh> cmt broadcast -local gmake (5) sh> cmt
broadcast -depth=<n> gmake (6) sh> cmt broadcast
-all_packages gmake (7)
```

According to the option, the loop will only operate onto:

1. the first package which path contains the string "Cm", and then all other reachable packages (in case other specifiers are used)
2. the packages which path contains the string "Cm"
3. the packages which path contains either the string "/Cm/" or the string "/CSet/"
4. the packages which path contains the string "Cm", but which does not contain the string "Cmo"
5. the packages at the same level as the current package
6. the packages at the same level as the current package or among the <n> first entries in the CMTPATH list
7. all the packages and versions currently available through the CMTPATH list

12.3.2 - cmt build <option>

All build commands are generally meant to change the current package (some file or set of files is generated). Therefore a check against conflicting locks (ie. a lock owned by another user) is performed by all these commands prior to execute it.

- [-nmake] constituent_makefile
<constituent-name>

This command is internally used by CMT in the standard Makefile.header fragment. It generates a specific makefile fragment (named *<constituent-name>.make*) which is used to re-build this fragment.

All such constituent fragments are automatically included from the main Makefile.

Although this command is meant to be used internally (and transparently) by CMT when the make command is run, a developer may find useful in very rare cases to manually re-generate the constituent fragment, using this command.

The -nmake option (which must precede the command) provides exactly the same features but in a Windows/nmake context. In this case, all generated makefiles are suffixed by .nmake instead of .make for Unix environments. The main makefile is expected to be named NMake and the standard header is named NMakefile.header

- **[-nmake] constituents_makefile**

This command is internally (and transparently) used by CMT in the standard Makefile.header fragment, and when the make command is run, to generate a specialized make fragment containing all "cmt build constituent_makefile" commands for a given package.

The -nmake option (which must precede the command) provides exactly the same feature but in a Windows/nmake context. In this case, all generated makefiles are suffixed by .nmake instead of .make for Unix environments. The main makefile is expected to be named NMake and the standard header is named NMakefile.header

- **dependencies**

This command is internally (and transparently) used by CMT from the constituent specific fragment, and when the make command is run, to generate a fragment containing the dependencies required by a source file.

This fragment contains a set of macro definitions (one per constituent source file), each containing the set of found dependencies.

- **library_links**

This command builds a local symbolic link towards all exported libraries from the used packages. A package exports its libraries through the `<package>_libraries` macro which should contain the list of constituent names corresponding to libraries that must be exported.

```
library Foo ... library Foo-utils ... ... macro
Foo_libraries "Foo Foo-utils" ...
```

The corresponding `cmt remove library_links` command will remove all these links.

- **make_setup**

This command is internally (and transparently) used by CMT from the standard `Makefile.header` fragment, and when the `make` command is run, to generate another fragment containing all platform (or tag) specific macro definitions.

One copy of this fragment (named `<tag>.make`) is created per flavour of tag used at build time. The tag considered in this operation is either the default tag value (obtained from the `CMTCONFIG` environment variable) or specified to the `make` command using the `-tag=<tag>` option)

This tag specific fragment represents the actual context that was considered at the most recent `make` activation. It is automatically rebuilt when one of the used requirements is modified.

- **msdev**

This command generates workspace (.dsw) and project (.dsp) files required for the MSDev tool.

- **os9_makefile**

This command generates external dedicated *makefile* fragments for each individual component of the package (ie. libraries or executable applications) to be used in OS9 context. It generates specific syntaxes for the OS9 operating systems.

The output of this tool is a set of files (named with the components' name and suffixed by .os9make) that are meant to be *included* within the main Makefile that the developer has to write anyhow.

The syntax of the cmt build os9_makefile utility is as follows :

```
sh> cmt build os9_makefile  
<package>
```

- **prototype <source-file-name>**

This command is internally (and transparently) used by CMT from the constituent specific fragment, and when the make command is run, to generate prototype header files from C source files.

The prototype header files (named <file-name>.ph) will contain prototype definitions for every global entry point defined in the corresponding C source file.

The effective activation of this feature is controled by the build strategy of CMT. The build strategy may be freely and globally overridden from any

requirements file, using the **build_strategy cmt** statement, providing either the **"prototypes"** or the **"no_prototypes"** values.

In addition, any constituent may locally override this strategy using the **"-prototypes"** or **"-no_prototypes"** modifiers.

- **readme**

This command generates a **README.html** file into the **cmt** branch of the referenced package. This html file will include

- a table containing URLs to equivalent pages for all used packages,
- a copy of the local **README** file (if it exists).

- **tag_makefile**

This command produces onto the standard output, the exhaustive list of all macros controled by CMT, ie. those defined in the requirements files as well as the standard macros internally built by CMT, taking into account all used packages.

12.3.3 - cmt check configuration

This command reads the hierarchy of requirements files referenced by a package, check them, and signals syntax errors, version conflicts or other configuration problems.

An empty output means that everything is fine.

12.3.4 - cmt check files <reference-file> <new-file>

This command compares the reference file to the new file, and substitutes the reference file by the new one if they are different.

If substitution is performed, a copy (with additional extension sav) is produced.

12.3.5 - cmt checkout ...

See the paragraph on how to use cvs together with CMT, and more specifically the details on checkout operations.

12.3.6 - cmt co ...

This is simply a short cut to the cmt checkout command.

12.3.7 - cmt cleanup [-csh|-sh]

This command generates (to the standard output) a set of shell commands (either for csh or sh shell families) meant to unset all environment variables specified in the requirements files of the used packages.

This command is internally used in the cleanup.[c]sh shell script, itself generated by the cmt config command.

12.3.8 - cmt config

This command (re-)generates the setup scripts and the minimal Makefile (when it does not exist yet or have been lost).

```
csh> cd ~/Packages/Foo/v1/cmt csh> cmt  
config
```

To be properly operated, one must *already* be in the cmt or mgr branch of a package (where the requirements file can be found).

This command also performs some cleanup operations (eg. removing all makefile fragments previously generated). Generally speaking, one may say that this command restores the CMT-related files to their original state (ie before any document generation)

The situations in which it is useful to run this command are:

- **When the setup or cleanup scripts have been lost**
 - **When the minimal Makefile have been lost**
 - **When the version of CMT is changed**
 - **After restoring a package from CVS**
 - **After having manually changed the directory structure of a package (using a manual copy operation, or renaming one of its parent directory, such as the version directory)**
-

12.3.9 - cmt create <package> <version> [<area>]

This command creates a new package or a new version of a package

```
csh> cmt create Foo
v1
```

or:

```
csh> cmt create Foo v1
~/dev
```

In the first mode (ie. without the *area* argument) the package will be created in the default path.

The second mode explicitly provides an alternate path.

A minimal configuration is installed for this new package:

- **An src and an cmt branch**
- **A very minimal requirements file**
- **The setup and cleanup shell scripts**
- **The minimal Makefile**

12.3.10 - cmt filter <in-file> <out-file>

This command reads <in-file>, substitutes all occurrences of macro references (taking either the form $\$(macro-name)$ or $\{\$macro-name\}$) by values deduced from corresponding macro specifications found in the requirements files, and writes the result into <out-file>.

This mechanism is widely internally used by CMT, especially for instanciating make fragments. Then, users may use it for any kind of document, including maual generation of MSDev configuration files, etc...

12.3.11 - cmt help

This command shows the list of options of the cmt driver.

12.3.12 - cmt lock

cmt lock [<package> <version> [<area>]]

This command tries to set a lock onto the current package (or onto the specified package). This consists in the following operations:

- 1. Check if a conflicting lock is already set onto this package (ie. a lock owned by another user).**
- 2. If not, then install a small text file named lock.cmt into the cmt/mgr branch of the package, containing the following text:**

```
locked by <user-name> date  
<now>
```

- 3. Run a shell command described in the macro named lock_command meant to install physical locks onto all files for this version of this package. A typical definition for this macro could be:**

```
macro lock_command "chmod -R a-w ../*" \ WIN32 "attrib  
/S /D +R ../*"
```

12.3.13 - `cmt remove <package> <version> [<area>]`

This command removes one version of the specified package. If the package does not contain a conflicting lock, and if the user is granted enough access rights to remove files, *all* files below the version directory will be definitively removed. Therefore this command should be used with as much care as possible.

The arguments needed to reach the package version to be removed are the same as the ones which had been used to create it.

If the removed version is the last version of this package, (and only if its directory is really empty) the package directory itself will be deleted.

12.3.14 - `cmt remove library_links`

This command removes symbolic links towards all imported libraries which had been installed using the `cmt build library_links` command. This command is generally transparently executed when one runs `gmake clean`

12.3.15 - `cmt run shell-command`

This command runs any shell command, in the context of the current package.

This may not appear to be very useful for the current package one is working on, but when combined with global options `-pack=package`, `-version=version`, `-path=access-path`, this gives a direct access to any

package context.

***12.3.16* - cmt setup [-csh|-sh|-bat]**

This command generates (to the standard output) a set of shell commands (either for csh, sh or bat shell families) meant to set all environment variables specified in the requirements files of the used packages.

This command is internally used in the setup.[c]sh or setup.bat shell script, itself generated by the cmt config command.

***12.3.17* - cmt show <option>**

- **author**
- **branches**
- **clients <package> [<version>]**

This command displays all packages that express an explicit use statement onto the specified package. If no version is specified on the argument list, then all uses of that package are displayed.

- **constituent_names**
- **constituents**
- **uses**
- **fragment <name>**

This command displays the actual location where the specified make fragment is currently found by CMT, taking into account possible overridden definitions.

- **fragments**
- **groups**

This command displays all groups possibly defined in constituents of the current package (using the `-group=<group-name>` option).

- **languages**
- **macro <name>**

This command displays a quite detailed explanation on the value assigned to the macro specified as the additional argument. It presents in particular each intermediate assignments made to this macro by the hierarchy of used statements, as well as the final result of these assignment operations.

By adding a `-tag=<tag>` option to this command, it is possible to simulate the behaviour of this command in another context, without actually going to a machine or an operating system where this configuration is defined.

- **macro_value <name>**

This command displays the raw value assigned to the macro specified as the additional argument. It only presents the final result of the assignment operations performed by used packages.

By adding a `-tag=<tag>` option to this command, it is possible to simulate the behaviour of this command in another context, without actually going to a machine or an operating system where this configuration is defined.

The typical usage of the `show macro_value` command is to get at the shell level (rather than within a Makefile) the value of a macro definition,

providing means of accessing them (quite similarly to an environment variable) :

```
csH> set compiler='cmt show macro_value cppcomp' csh>
${compiler} ....
```

- **macros**

This command extracts from the requirements file(s) the complete set of macro definitions, selects the appropriate *tag* definition (or uses the one provided in the `-tag=<tag>` option) and displays the effective macro values corresponding to this tag.

This command is typically used to show the effective list of macros used when running make and can be also used to build, as an argument list, the make command as follows :

```
csH> eval make `cmt show
macros`
```

This use of `cmt show macros` is directly installed within the default target provided in the standard `Makefile.header` file. Therefore if this file is included into the package's `Makefile`, macro definitions provided in the requirements files (the one of the currently built package as well as the ones of the used packages) will be expanded and provided as arguments to `make`.

By adding a `-tag=<tag>` option to this command, it is possible to simulate the behaviour of this command in another context, without actually going to a machine or an operating system where this configuration is defined.

- **manager**

- **packages**

This command displays all packages (and all versions of them) currently reachable through the current access path definition (which can be displayed using the `cmt show path` command).

- **path**

This command displays the complete and effective access path currently defined using any possible alternate way.

- **pattern <name>**

This command displays how and where the specified pattern is defined, and which packages do apply it.

- **patterns**

This command displays all pattern definitions.

- **pwd**

This command displays a filtered version of the standard `pwd` unix command. The applied filter takes into account the set of aliases installed in the standard configuration file located in `${CMTROOT}/mgr/cmt_mount_filter`.

This configuration file contains a set of path aliases (one per line) each proposing a translation for non-portable file paths (imposed by mount constraints on some contexts).

- **set_value <name>**

- **set <name>**

- **sets**

- **strategies**

- tags
- uses

This command displays a quite comprehensive and detailed explanation of the hierarchy of use statements, with the effective selection made between possibly compatible versions.

```
# use CMT v1r10 /lal # use Cm v7r5 # use CSet v2r5 # #
Selection : use CSet v2r5 /lal use Cm v7r5 /lal use CMT
v1r10 /lal
```

The -quiet option may be used to remove from the output, the comments (beginning with the # character), so as to display a simple list of used packages, starting from the deepest uses.

- version

This command displays the version tag of the current package.

- versions <name>

This command displays the reachable versions of the specified package, looking at the current access paths.

12.3.18 - cmt system

This command displays the current value assigned by default to the CMTCONFIG environment variable.

12.3.19 - cmt unlock

cmt unlock [<package> <version> [<area>]]

This command tries to remove a lock from the current package (or from the specified package).

This consists in the following operations:

1. Check if a conflicting lock is already set onto this package (ie. a lock owned by another user).
2. If not, then remove the text file named lock.cmt from the cmt/mgr branch of the package.
3. Run a shell command described in the macro named unlock_command meant to remove physical locks from all files for this version of this package. A typical definition for this macro could be:

```
macro unlock_command "chmod -R g+w ../*" \ WIN32
"attrib /S /D -R ../*"
```

12.3.20 - cmt version

This command shows the current version of CMT, including (if applicable) the actual patch level. This always corresponds to the corresponding CVS tag assigned to CMT sources.

12.3.21 - cmt cvstags <module>

(see the section on *how to use CVStogether with CMT* for more details on this command)

12.3.22 - cmt cvsbranches <module>

12.3.23 - cmt cvssubpackages <module>

12.4 - The setup and cleanup scripts

They are produced by the cmt config command and their contents is built according to the specifications stored in the requirements file.

One flavour of these scripts is generated per shell family (csh, sh and bat), yielding the following scripts :

```
setup.csh setup.sh setup.bat cleanup.csh  
cleanup.sh
```

The main sections installed within a setup script are :

- 1. Connection to the current version of the CMT package.**
- 2. Setting the set of user defined public variables specified in the requirements file (including those defined by all used packages). This is achieved by running the cmt setup utility into a temporary file and running this temporary file.**
- 3. Activation of the user defined setup and cleanup scripts (those specified using the setup_script and cleanup_script statements).**

It should be noted that these setup scripts do *not* contain machine specific information (due to the online use of the cmt setup command). Therefore, it is perfectly possible to use the same setup script on various platforms (as soon as they share the directories) and this also shows that the configuration operation (the cmt config command) is required only once for a set of multiple platforms sharing a development area.

12.5 - cmt build prototype

This command is only provided for development of C modules. It generates a C header file containing the set of prototype statements for all public functions of a given module. Its output is a file with the same name as the input source (given as the argument) and suffixed

with .ph.

The generated prototype header file is meant to be included wherever it is needed (in the module file itself for instance).

A typical example of the use of cmt build prototype could be :

```
ssh> cd ../src ssh> cmt build prototype FooA.c Building
FooA.ph
```

Running cmt build prototype will only produce a new prototype header file if the output is actually different from the existing one (if it exists) in order to avoid confusing *make* checks.

The effective use of this facility (which may not be appropriate in all projects) is controlled by one option of the build strategy, which can take one of the two values:

```
build_strategy prototypes build_strategy
no_prototypes
```

In addition to this global strategy specification, each application or library may individually override it using the *-prototypes* or *-no_prototypes* options.

Lastly, the actual behaviour of the prototype generator is defined in the standard make macro *build_prototype* (which default to call the *cmt build prototype* command, allowing a user defined behaviours to this feature)

13 - Using cvs together with CMT

Nothing special is apriori required by CMT with respect to the use of CVS. Nevertheless, one may advertize some well

tested conventions and practices which turned out to ensure a good level of consistency between the two utilities.

Although none of these are required, the `cmt general` command provides a few utilities so as to simplify the use of these practices. It should be noted that the added features provided by `cmt` rely on the possibility to *query* CVS about the existing CMT packages and the possible tags setup for these packages. CVS does not by default permit such query operations (since they require to scan the physical CVS repository). Therefore CMT provides a hook to CVS (based upon standard CVS features - not patches) for this. This hook can be installed by the following procedure (see sections below for more details):

```
sh> (cd ${CMTROOT}/mgr; gmake
installcvs)
```

13.1 - Importing a package into a cvs repository

Generally, everything composing a package (below the *version* directory and besides the *binary* directories) is relevant to be imported. Then choosing a *cvs module* name is generally done on the basis of the package name. Taking the previous examples, one could import the Foo package as follows :

```
csh> cd ...../Foo/v1 csh> cvs import -m "First import" -I
alpha -I hp9000s700 Foo LAL v1
```

In this example,

- we have ignored the currently existing binary directories (here `alpha` and `hp9000s700`)
- the `cvs` module name is identical to the package name

(Foo)

- **the original symbolic insertion tag is identical to the version identifier (v1)**

The choice of the module name can generally be identical to the package name. However, some site specific management issues may lead to different choices (typically, a top directory where groups of packages are gathered may be inserted).

Conversely, using symbolic tags identical to version identifiers appears to be a very good practice. The only constraint induced by cvs is that the symbolic tags may not contain *dot* characters ('.'), whereas no restriction exist from CMT itself. Thus version identifiers like v3r2 will be preferred to the v3.2 form.

13.2 - Checking a package out from a cvs repository

Assuming the previous conventions on module name and version identifier have been selected when importing a package, the following operations will naturally intervene when one need to check a package out (typically to work on it or to install it on some platform) :

```

csh> cd <some root> (1) csh> mkdir Foo (2) csh> cd Foo csh>
cvs checkout -d v1 Foo (3) csh> cd v1/cmt csh> cmt config
(4) csh> source setup.csh (5) csh> [g]make (6)

```

1. one always have to select a root directory where to settle down this copy of the extracted package. This may either be the so-called *default root* or any other appropriate directory. In both cases, the next `cmt config` operation will automatically take care of this effective location.
2. creating a base directory with the package name is mandatory here, and is *not* taken into account by `cvs` during the *checkout* operation since one wants to insert the *version* branch in between.
3. the package is checked out into a directory named with the expected version identifier exactly corresponding to the version currently stored in the `cvs` repository.
4. then using the `cmt config` command (from the `cmt` branch) will update the setup scripts against the `requirements` file and the effective current package location.
5. using this updated version of the setup script provides the appropriate set of environment variables
6. lastly, rebuilding the entire package is possible simply using the `[g]make` command.

The actions described just above (from number 2 to number 4 included) can also be performed using the `cmt checkout` command.

```

> cd <some work area> > cmt checkout [modifier ...]
<package> ... modifier : -l Do not process used packages
(default). -R Process used packages recursively. -r rev
Check out version tag. (is sticky) -d dir Check out into dir
instead of module name. -o offset Offset in the CVS
repository -n Simulation mode on -v Verbose mode on -help
Print this help

```

Thus the previous example would become:

```

csh> cd <some root> csh> cmt checkout Foo csh> cd Foo/v1/cmt
csh> source setup.csh csh> [g]make

```

13.3 - Querying CVS about some important infos

It is possible, using the commands :

- `cmt cvstags <module>`
- `cmt cvsbranches <module>`

- **cmt cvssubpackages <module>**
to query the CVS repository about the existing tags installed onto a given CVS module, the subdirectories and the subpackages (in the CMT meaning, i.e. when a requirements file exists).

```
> cmt cvstags Cm v7r6 v7r5 v7r4 v7r3 v7r1 v7 > cmt cvstags  
Co v3r7 v3r6 v3
```

One should notice here that the `cvstags` command can give informations about any type of module, even if it is not managed in the CMT environment.

However, in order to let this mechanism operate, it is required to install some elements into the physical CVS repository (*which may require some access rights into it*). This installation procedure (to be done only once in the life of the repository) can be achieved through the following command:

```
sh> (cd ${CMTROOT}/mgr; gmake  
installcvs)
```

However, the details of the procedure is listed below (this section is preferably reserved for system managers and can easily be skipped by standard users):

1. copy the `cmt_buildcvsinfos2.sh` shell script into the management directory `${CVSROOT}/CVSROOT` as follows :

```
sh> cp ${CMTROOT}/mgr/cmt_buildcvsinfos2.sh  
${CVSROOT}/CVSROOT
```

2. install one special statement in the `loginfo` administrative file as follows :

```
sh> cd ... sh> cvs checkout CVSROOT sh> cd CVSROOT sh> vi
loginfo ... .cmtcvsinfos
$CVSROOT/CVSROOT/cmt_buildcvsinfos2.sh sh> cvs commit -m
"set up commitinfo for CMT"
```

13.4 - Working on a package, creating a new release

This section presents the way to instantiate a new release of a given package, which happens when the foreseen modifications will yield additions or changes to the application programming interface of the package.

Then the version tag is supposed to be moved forward, either increasing its minor identifier (in case of simple additions) or its major identifier (in case of changes).

The following actions should be undertaken then :

- 1. understand what is the latest version tag (typically by using the `cmt cvstags` command). Let's call it `old-tag`.**
- 2. select (according to the foreseen amount of changes) what will be the next version tag. Let's call it `new-tag`.**
- 3. check the most recent version of the package in your development area :**

```
sh> cd <development area> sh> cvs checkout -d <new-tag>
<package>
```

- 4. configure this new release, and rebuild it :**

```
sh> cd <new-tag>/cmt sh> cmt config sh> source setup.csh
sh> [g]make
```

13.5 - Getting a particular tagged version out of cvs

The previous example presented the standard case where one gets the *most recent* version of a given package. The procedure is only slightly modified when one wants to extract a previously tagged version. Let's imagine that the Foo package has evolved by several iterations, leading to several tagged releases in the cvs repository (say v2 and v3). If the v2 release is to be used (e.g. for understanding and fixing a problem discovered in the running version) one will operate as follows :

```
ssh> cd <some root> ssh> mkdir Foo ssh> cd Foo ssh> cvs
checkout -d v2 -r v2 Foo ssh> cd v2/cmt ssh> cmt config ssh>
source setup.csh ssh> make
```

14 - Interfacing an external package with CMT

Very often, external packages (typically commercial products, or third party software) are to be used by packages developed in the context of the CMT environment. Although this can obviously be done simply by specifying compiler or linker options internally to the client packages, it can be quite powerful to interface these so-called *external* packages to CMT by defining a *glue* package, where configuration specifications for this external package are detailed.

Using this approach, one may :

- provide a *nickname* for this external package,
- adapt the version tag convention consistently to the project, hiding the version tag specificities of eg. commercial packages.
- provide compiler options using the standard make

macros `<package>_cflags`, `<package>_cppflags` or `<package>_fflags`,

- specify a set of search paths for the include files, using the `include_dirs` statement,
- provide linker options using the the standard make macros `<package>_linkopts`

Let's consider the example of the OPACS package. This package is provided outside of the CMT environment. Providing a directory tree following the CMT conventions (ie. a branch named after the version identifier, then an cmt branch) then a requirements file, containing (among other statements not shown for the sake of clarity) :

```
package OPACS include_dirs ${Wo_root}/include
${Co_root}/include ${Xx_root}/include \ ${Ho_root}/include
${Go_root}/include ${Xo_root}/include public macro
OPACS_cflags "-DHAS_XO -DHAS_XM" macro OPACS_cppflags "
$(OPACS_cflags) " macro OPACS_linkopts "$(Wo_linkopts)
$(Xo_linkopts) $(Go_linkopts) \ $(Glo_linkopts) $(Xx_linkopts)
$(Ho_linkopts) $(Html0_linkopts) \ $(W3o_linkopts)
$(Co_linkopts) $(X_linkopts)"
```

Then every package or application, client of this OPACS package would have just to provide a use statement like :

```
use OPACS
v3
```

This procedure gives the complete benefit of the use relationships between packages (a client application transparently inherits all configuration specifications) while keeping unchanged the original referenced package, allowing to apply this approach even to commercial products so that they may be integrated in resource usage surveys similarly to local packages.

15 - Installing CMT for the first time

These sections are of interest only if CMT is not yet installed on your site, or if you need a private installation.

The first question you need to answer is the location where to install CMT. This location is typically a disk area where most of packages managed in your project will be located.

Then, you have to fetch the distribution kit from the Web at <http://www.lal.in2p3.fr/SI/CMT/CMT.htm>. You must get at least the primary distribution kit containing the basic configuration information and the CMT sources. This operation results in a set of directories hanging below the CMT root and the version directory. The src branch contains the sources of CMT, the fragments branch contains the makefile fragments and the mgr branch contains the scripts needed to build or operate CMT.

15.1 - Installing CMT on your Unix site

The very first operation after downloading CMT consists in running the INSTALL shell script. This will build the setup scripts required by any CMT user.

Then you may either decide to build CMT by yourself or fetch a pre-built binary from the same Web location. The prebuilt binary versions may not exist for the actual platform you are working on. You will see on the distribution page the precise configurations used for building those binaries.

In case you have to build CMT yourself, you need a C++ compiler capable of handling templates (although the

support for STL is not required). There is a Makefile provided in the distribution kit which takes g++ by default as the compiler. If you need a specific C++ compiler you will override the cpp macro as follows:

```
sh> gmake  
cpp=CC
```

The cppflags macro can also be used to override the behaviour of the compilation.

Another important concern is the way CMT will identify the platform. CMT builds a configuration tag per each type of platform, and uses this tag for naming the directory where all binary files will be stored. As such this tag has to be defined prior to even build CMT itself.

CMT builds the default configuration by running the cmt_system.sh script found in the mgr branch of CMT. Run it manually to see what is the default value provided by this script. You might consider changing its algorithm for your own convenience.

A distribution kit may be obtained at the following URL :

```
http://www.la1.in2p3.fr/SI/CMT/CMT.htm
```

Once the tar file has been downloaded, the following operations must be achieved :

1. Select a root directory where to install CMT.
Typically this will correspond to a development area or a public distribution area.
2. Import the distribution kit mentioned above.
3. Uncompress and untar it.

4. Configure CMT.

5. CMT is ready to be used for developing packages.

A typical corresponding session could look like :

```
ssh> cd /Packages ssh> <get the tar file from the Web> ssh>
uncompress CMTv1r10.tar.Z ssh> tar xvf CMTv1r10.tar ssh> cd
CMT/v1r10/mgr ssh> ./INSTALL ssh> source setup.csh ssh>
gmake
```

15.2 - Installing CMT on a Windows or Windows NT site

You first have to fetch the distribution kit from the Web at <http://www.lal.in2p3.fr/SI/CMT/CMT.htm>. You must get at least the primary distribution kit containing the basic configuration information and the CMT sources. This operation results in a set of directories hanging below the CMT root and the version directory. The binary kit provided for Windows environments will generally fit your needs.

You should consider getting the pre-compiled (for a Windows environment) applications, and especially the `..\VisualC\install.exe` application, which interactively configures the registry entries as described in the next paragraph.

The next operation consists in defining a few registries (typically using the standard RegEdit facility or the `install.exe` special application):

- `HKEY_LOCAL_MACHINE/Software/CMT/root` will contain the root directory where CMT is installed (eg. "e:").
- `HKEY_LOCAL_MACHINE/Software/CMT/version` will contain the current version tag of CMT ("v1r10")

for this version).

- **HKEY_LOCAL_MACHINE/Software/CMT/path/** may optionally contain a set of text values corresponding to the different package global access paths.
- **HKEY_LOCAL_MACHINE/Software/CMT/site** will contain the conventional site name.
- **HKEY_CURRENT_USER/Software/CMT/path/** may contain a set of text of text values corresponding to the different package private access paths.

CMT can also be configured to run on DOS-based environments using the nmake facility. In this case, the installation procedure is very similar to the Unix one:

A typical corresponding session could look like :

```
dos> cd Packages dos> <get the tar file from the Web> dos>  
cd CMT\v1r10\mgr dos> call INSTALL dos> call setup.bat dos>  
nmake /f nmake
```

16 - Differences with previous versions of CMT

16.1 - Converting a package that was managed with previous versions of CMT (or methods)

The primary source of information handled by CMT, i.e. the syntax - and semantics - of the requirements file is supposed to be maintained backward compatible across previous versions of CMT. Therefore we expect that the effects of using a new version of CMT to a package already managed by previous versions of CMT, will remain limited.

Generally, it is enough to just *re-configure* the package, using the well known command

```
sh> cmt
config
```

This will result in re-generating the setup scripts, and verifying Makefile. It will also remove all existing makefile fragments previously generated by CMT. One should also check that a proper CMT-compliant Makefile should contain at least the two following lines:

```
include ${CMTROOT}/src/Makefile.header include
${CMTROOT}/src/constituents.make
```

These two lines are the only required lines to be present in an operational Makefile. However, the user is entirely free to install additional make statements at any location for his/her own purpose.

No further operation is then needed. All other makefile fragments will be automatically re-generated at make time.

Sometimes, it might also be useful to clean the binary directories and rebuild the package as follows:

```
sh> gmake clean sh>
gmake
```

Lastly, it's often useful to broadcast these actions (and primarily the cmt config action) towards all used packages at once. This of course can easily be done through the cmt broadcast command as follows:

```
sh> cmt broadcast cmt
config
```

16.2 - Operations in a Windows environment

A graphical and interactive application (cmtw) is provided on Windows (95/98/NT) environments. This application let the developer browse the package directories, select any version of any package. Its configuration is shown, and interactive edition is possible on its requirements file. A few operations are also possible, such as the generation of MSDev configuration files, so as to directly work with packages managed by CMT within the MSDev development environment. Currently this support is still evolving and user might see limitations in the dialog between CMT and MSDev (only the constituent definitions - applications and libraries - and the use mechanism - package relationships - are understood in the context of MSDev). Users of these new facilities are kindly invited to send their comments, bug observations, suggestions or even contributions to the author.

17 - Appendices

17.1 - Standard make targets predefined in CMT

These targets can always be listed through the following command :

```
sh> gmake  
help
```

<i>target</i>	<i>usage</i>
help	Get the list of possible make target for this package.
all	build all components of this package.
clean	remove everything that can be rebuilt by make
configclean	remove all intermediate makefile fragments
check	run all applications defined with the -check option
<i>component-name</i>	only build this particular component (as opposed to the all target that tries to build all components of this package)
<i>group-name</i>	build all constituents belonging to this group (ie. those defined using the same -group= option)

These targets have to be specified as follows :

```
sh> gmake clean sh> gmake
Foo
```

17.2 - Standard macros predefined in CMT

17.2.1 - Structural macros

These macros describe the structural conventions followed by CMT. They receive a conventional default value from the CMT requirements file. However, they can be overridden in any package for its own needs.

<i>macro</i>	<i>usage</i>	<i>default value</i>
tag	gives the binary tag	<code>\${CMTCONFIG}</code>
src	the src branch	<code>../src/</code>
inc	the include branch	<code>../src/</code>
mgr	the cmt or mgr branch	<code>../cmt/ or ../mgr/</code>
bin	the branch for binaries	<code>../\${CMTCONFIG}/</code>
javabin	the branch for java classes	<code>../classes/</code>
doc	the doc branch	<code>../doc/</code>

17.2.2 - Language related macros

These macros are purely conventional. They are expected in the various make fragments available from CMT itself for providing the various building actions.

During the mechanism of new language declaration and definition available in the CMT syntax, developers are expected to define similar conventions for corresponding actions.

Their default values are originally defined inside the requirements file of the CMT package itself but can be *redefined* by providing a new definition in the package's requirements file using the macro statement. The original definition can be *completed* using the macro `macro_append` or `macro_prepend` statements.

<code>cc</code>	The C compiler	<code>cc</code>
<code>ccomp</code>	The C compiling command	<code>\$(cc) -c -I\$(inc) \$(includes) \$(cflags)</code>
<code>clink</code>	The C linking command	<code>\$(cc) \$(clinkflags)</code>
<code>cflags</code>	The C compilation flags	<i>none</i>
<code>pp_cflags</code>	The preprocessor flags for C	<i>none</i>
<code>clinkflags</code>	The C link flags	<i>none</i>

cpp	The C++ compiler	g++
cppomp	The C++ compiling command	\$(cpp) -c -I\$(inc) \$(includes) \$(cppflags)
cpplink	The C++ linking command	\$(cpp) \$(cpplinkflags)
cppflags	The C++ compilation flags	<i>none</i>
pp_cppflags	The preprocessor flags for C++	<i>none</i>
cpplinkflags	The C++ link flags	<i>none</i>

for	The Fortran compiler	f77
fcomp	The Fortran compiling command	\$(for) -c -I\$(inc) \$(includes) \$(fflags)
fink	The Fortran linking command	\$(for) \$(clinkflags)
fflags	The Fortran compilation flags	<i>none</i>
pp_fflags	The preprocessor flags for fortran	<i>none</i>
finkflags	The Fortran link flags	<i>none</i>
ppcmd	The include file command for Fortran	-I

javacomp	The java compiling command	javac
jar	The java archiver command	jar

lex	The Lex command	lex \$(lexflags)
lexflags	The Lex flags	<i>none</i>
yacc	The Yacc command	yacc \$(yaccflags)
yaccflags	The Yacc flags	<i>none</i>
ar	The archive command	ar -clr
ranlib	The ranlib command	ranlib

17.2.3 - Package customizing macros

These macros do not receive default values. They are all prefixed by the package name. They are meant to provide specific variant to the corresponding generic language related macros.

They are automatically and by default concatenated by CMT to fill in the corresponding global *use* macros (see appendix on generated macros). However, this concatenation mechanism is discarded when the `-no_auto_imports` option is specified in the corresponding use statement.

<code><package>_cflags</code>	specific C flags
<code><package>_pp_cflags</code>	specific C preprocessor flags
<code><package>_cppflags</code>	specific C++ flags
<code><package>_pp_cppflags</code>	specific C++ preprocessor flags
<code><package>_fflags</code>	specific Fortran flags
<code><package>_pp_fflags</code>	specific Fortran preprocessor flags
<code><package>_libraries</code>	gives the (space separated) list of library names exported by this package. This list is typically used in the <code>cmt build library_links</code> command.
<code><package>_linkopts</code>	<p>provide the linker options required by any application willing to access the different libraries offered by the package. This may include support for several libraries per package.</p> <p>A typical example of how to define such a macro could be :</p> <pre>macro Cm_linkopts "-L\$(CMROOT)/\$(Cm_tag) -lCm -lm"</pre>
<code><package>_stamps</code>	<p>may contain a list of <i>stamp</i> file names (or make targets). Whenever a library is modified, one dedicated stamp file is re-created, simply to mark the reconstruction date. The name of this stamp file is conventionally <code><library>.stamp</code>. Thus, a typical definition for this macro could be :</p> <pre>macro Cm_stamps "\$(Cm_root)/\$(Cm_tag)/Cm.stamp"</pre> <p>Then, these stamp file references are accumulated into the standard macro named <code>use_stamps</code> which is always installed within the dependency list for applications, so that whenever one of the libraries used from the hierarchy of used packages changes, the application will be automatically rebuilt.</p>

17.2.4 - Constituent specific customizing macros

These macros do not receive any default values (ie they are empty by default). They are meant to provide for each constituent, specific variants to the corresponding generic language related macros.

By convention, they are all prefixed by the

constituent name. But macros used for defining compiler options are in addition prefixed by the constituent category (either lib_, app_ or doc_

They are used in the various make fragments for fine customization of the build command parameters.

<code><category>_<constituent>_cflags</code>	specific C flags
<code><category>_<constituent>_pp_cflags</code>	specific C preprocessor flags
<code><category>_<constituent>_cppflags</code>	specific C++ flags
<code><category>_<constituent>_pp_cppflags</code>	specific C++ preprocessor flags
<code><category>_<constituent>_fflags</code>	specific Fortran flags
<code><category>_<constituent>_pp_fflags</code>	specific Fortran preprocessor flags
<code><constituent>linkopts</code>	provides additional linker options to the application. It is complementary to - and should not be confused with - the <code><package>_linkopts</code> macro, which provides exported linker options required by clients packages to use the package libraries.
<code><constituent>_shlibflags</code>	provides additional linker options used when building a shared library. Generally, a simple shared library does not need any external reference to be resolved at build time (it is in this case supposed to get its unresolved references from other shared libraries). However, (typically when one builds a dynamic loading capable component) it might be desired to statically link it with other libraries (making them somewhat private).
<code><constituent>_dependencies</code>	provides user defined dependency specifications for each constituent. The typical use of this macro is fill it with the name of the list of some other constituents which <i>have</i> to be rebuilt first (since each constituent is associated with a target with the same name). This is especially needed when one want to use the parallel gmake (ie. the -j option of gmake).

17.2.5 - Source specific customizing macros

These macros do not receive any default values (ie they are empty by default). They are meant to provide for each source file, specific variants to the corresponding generic language related macros.

By convention, they are all prefixed by the source file name followed by the source file suffix (either `_c`, `_cxx`, `_f`, etc.)

They are used in the various make fragments for fine customization of the build command parameters.

<code><constituent>_<suffix>_cflags</code>	specific C flags
<code><constituent>_<suffix>_cppflags</code>	specific C++ flags
<code><constituent>_<suffix>_fflags</code>	specific Fortran flags

17.2.6 - Generated macros

These macros are automatically *generated* when make is run.

The first set of them provide constant values corresponding to CMT based information. They are not meant to be overridden by the user, since they serve as a communication mean between CMT and the user.

<code><PACKAGE>ROOT</code>	The access path of the package (including the version branch)
<code><package>_root</code>	The access path of the package (including the version branch). This macro is very similar to the <code><PACKAGE>ROOT</code> macro except that it tries to use a relative path instead of an absolute one.
<code><PACKAGE>VERSION</code>	The used version of the package
<code>PACKAGE_ROOT</code>	The access path of the current package (including the version branch)
<code>package</code>	The name of the current package
<code>version</code>	The version tag of the current package

The second set is deduced from the context and from the requirements file of the package. They can be overridden by the user so as to customize the CMT behaviour.

<code><package>_tag</code>	The specific configuration tag for the package. By default it is set to <code>\$(tag)</code> but can be freely overridden
<code>constituents</code>	The ordered set of constituents declared without any group option
<code><group-name>_constituents</code>	The ordered set of all constituents declared using a <code>group=<group-name></code> option

The third set of generated macros are the *global use macros*. They correspond to the concatenation of the corresponding package specific customizing options that can be deduced from the ordered set of *use* statements found in the requirements file (taking into account the complete hierarchy of used packages with the exception of those specified with the `-no_auto_imports` option in their use statement) :

use_cflags	C compiler flags
use_pp_cflags	Preprocessor flags for the C language
use_cppflags	C++ compiler flags
use_pp_cppflags	Preprocessor flags for the C++ language
use_fflags	Fortran compiler flags
use_pp_fflags	Preprocessor flags for the Fortran language
use_libraries	List of library names
use_linkopts	Linker options
use_stamps	Dependency stamps
use_requirements	The set of used requirements
use_includes	The set of include search paths options for the preprocessor from the used packages
use_fincludes	The set of include search paths options for the fortran preprocessor from the used packages
includes	The overall set of include search paths for the preprocessor
fincludes	The overall set of include search paths options for the fortran preprocessor

17.2.7 - Utility macros

These macros are used to specify the behaviour of various actions in CMT.

X11_cflags	compilation flags for X11
Xm_cflags	compilation flags for Motif
X_linkopts	Link options for XWindows (and Motif)
make_shlib	The command used to generate the shared library from the static one
shlibsuffix	The system dependent suffix for shared libraries
shlibbuilder	The loader used to build the shared library
shlibflags	The additional options given to the shared library builder
symlink	The command used to install a symbolic link
	The command used to remove a symbolic link
build_prototype	The command used to generate the C prototype header file (default to the internal cmt dedicated command)
build_dependencies	The command used to generate dependencies (default to the internal cmt dedicated command)
lock_command	The command used to physically lock a package
unlock_command	The command used to physically unlock a package
make_hosts	The list of remote host names which exactly require the make command
gmake_hosts	The list of remote host names which exactly require the gmake command

17.3 - Standard templates for makefile fragments

<i>template name</i>	<i>usage</i>	<i>used in fragment</i>
ADDINCLUDE	additional include path	<language> java
CONSTITUENT	name of the constituent	<language> java jar make_header jar_header java_header library_header application_header protos_header library_no_share library application dependencies cleanup_header cleanup_library cleanup_application check_application document_header <document> trailer dsw_all_project_dependency dsw_project dsp_library_header dsp_shared_library_header dsp_windows_header dsp_application_header dsp_trailer constituent check_application_header
DATE	now	make_header

FILENAME	file name without path	buildproto <i><language></i> <i><document></i>
FILEPATH	file path	buildproto <i><language></i> <i><document></i>
FILESUFFIX	file suffix (without dot)	<i><language></i>
FILESUFFIX	file suffix (with dot)	<i><document></i>
FULLNAME	complete file path and name	<i><language></i> cleanup <i><document></i> dsp_contents
GROUP	group name	constituents_header
LINE	source files	<i><language></i> dependencies constituent
LINKMACRO	link macro	application
NAME	file name without path and suffix	buildproto <i><language></i> java <i><document></i>
OBJS	object files	jar_header java_header jar library_no_share library application cleanup_java document_header trailer
OUTPUTNAME	output file name	java
PACKAGE	current package name	<i><language></i> dsw_header dsw_all_project dsw_all_project_trailer dsw_trailer dsp_all make_setup_header make_setup readme_header readme readme_use readme_trailer
PACKAGEPATH	current package location	readme_use
PROTOSTAMPS	prototype stamp files	protos_header
PROTOTARGET	prototype target name	library_header application_header
SUFFIX	document suffix	<i><document></i>

TITLE	title for make header	make_header
USER	user name	make_header
VERSION	current package version tag	readme_header readme readme_use

17.4 - Makefile generation sequences

This section describes the various makefile generation sequences provided by CMT. Each sequence description shows the precise set of make fragments used during the operation.

<i>Generated makefile</i>	<i>description</i>	<i>used make fragments</i>
setup.make	Configuration files for make	<ol style="list-style-type: none"> 1. make_setup_header 2. make_setup
constituents.make	the main entry point point for all constituent targets	<ol style="list-style-type: none"> 1. constituents_header 2. constituent 3. check_application_header
<constituent>.make	application or library specific make fragment	<ol style="list-style-type: none"> 1. make_header 2. java_header jar_header library_header application_header 3. protos_header 4. buildproto 5. jar library library_no_share application 6. dependencies 7. <language> <language>_library java 8. cleanup_header 9. cleanup 10. cleanup_application 11. cleanup_objects 12. cleanup_java 13. cleanup_library 14. check_application

<code><constituent>.make</code>	document specific make fragment	<ol style="list-style-type: none"> 1. <code>make_header</code> 2. <code>document_header</code> 3. <code>dependencies</code> 4. <code><document></code> 5. <code><document-trailer></code> 6. <code>cleanup_header</code>
<code><package>.dsw</code>	Visual workspace configuration files	<ol style="list-style-type: none"> 1. <code>dsw_header</code> 2. <code>dsw_all_project_header</code> 3. <code>dsw_all_project_dependency</code> 4. <code>dsw_all_project_trailer</code> 5. <code>dsw_project</code> 6. <code>dsw_trailer</code> 7. <code>dsp_all</code>
<code><constituent>.dsp</code>	Visual project configuration files	<ol style="list-style-type: none"> 1. <code>dsp_library_header</code> <code>dsp_shared_library_header</code> <code>dsp_windows_header</code> <code>dsp_application_header</code> 2. <code>dsp_contents</code> 3. <code>dsp_trailer</code>
README	.	<ol style="list-style-type: none"> 1. <code>readme_header</code> 2. <code>readme</code> 3. <code>readme_use</code> 4. <code>readme_trailer</code>

17.5 - The complete requirements syntax

The syntax of specification statements that can be installed in a requirements file are :

cmt-statement : *alias*
| *application*
| *apply_pattern*
| *author*
| *branches*
| *build_strategy*
| *cleanup_script*
| *document*
| *ignore_pattern*

| *include_dirs*
| *include_path*
| *language*
| *library*
| *macro*
| *macro_append*
| *macro_prepend*
| *macro_remove*
| *macro_remove_all*
| *make_fragment*
| *manager*
| *package*
| *path*
| *path_append*
| *path_prepend*
| *path_remove*
| *pattern*
| *private*
| *public*
| *set*
| *set_append*
| *set_prepend*
| *set_remove*
| *setup_script*
| *tag*
| *use*
| *version*
| *version_strategy*

alias : *alias alias-name value [tag value ...]*

application : **application** *application-name* [*constituent-option* ...] *source*
 ...

constituent-option : **-OS9**
 | **-windows**
 | **-no_share**
 | **-no_static**
 | **-prototypes**
 | **-no_prototypes**
 | **-check**
 | **-group=group-name**
 | **-suffix=output-suffix**
 | **-import=package-name**
 | *variable-name=variable-value*

source : [**-s=new-search-path**] *file-name*

apply_pattern : **apply_pattern** *pattern-name* [*template-name=value* ...]

author : **author** *author-name*

branches : **branches** *branch-name* ...

build_strategy : **build_strategy** *build-strategy-name*

build-strategy-name : **prototypes**
 | **no_prototypes**
 | **keep_makefiles**
 | **rebuild_makefiles**

cleanup_script : **cleanup_script** *script-name*

document : **document** *document-name* [*constituent-option* ...] *source* ...

ignore_pattern : **ignore_pattern** *pattern-name*

include_dirs : **include_dirs** *search-path* ...

include_path : **include_path** *search-path*

language : **language** *language-name* [*language-option* ...]

language-option : **-suffix=suffix**
 | **-linker=linker-command**
 | **-prototypes**

		-preprocessor_command= <i>preprocessor-command</i>
		-fragment= <i>fragment</i>
		-output_suffix= <i>output_suffix</i>
		-extra_output_suffix= <i>extra_output_suffix</i>
<i>library</i>	:	library <i>library-name</i> [<u><i>constituent-option</i></u>] <i>source ...</i>
<i>macro</i>	:	macro <i>macro-name value</i> [<i>tag value ...</i>]
<i>macro_append</i>	:	macro_append <i>macro-name value</i> [<i>tag value ...</i>]
<i>macro_prepend</i>	:	macro_prepend <i>macro-name value</i> [<i>tag value ...</i>]
<i>macro_remove</i>	:	macro_remove <i>macro-name value</i> [<i>tag value ...</i>]
<i>macro_remove_all</i>	:	macro_remove_all <i>macro-name value</i> [<i>tag value ...</i>]
<i>make_fragment</i>	:	make_fragment <i>fragment-name</i> <u><i>fragment-option</i></u>
<i>fragment-option</i>	:	-suffix= <i>suffix</i>
		-dependencies
		-header= <i>fragment</i>
		-trailer= <i>fragment</i>
<i>manager</i>	:	manager <i>manager-name</i>
<i>package</i>	:	package <i>package-name</i>
<i>path</i>	:	path <i>path-name value</i> [<i>tag value ...</i>]
<i>path_append</i>	:	path_append <i>path-name value</i> [<i>tag value ...</i>]
<i>path_prepend</i>	:	path_prepend <i>path-name value</i> [<i>tag value ...</i>]
<i>path_remove</i>	:	path_remove <i>path-name value</i> [<i>tag value ...</i>]
<i>pattern</i>	:	pattern [-global] <i>pattern-name</i> <u><i>cmt-statement</i></u> [; <u><i>cmt-statement ...</i></u>]
<i>private</i>	:	private
<i>public</i>	:	public
<i>set</i>	:	set <i>set-name value</i> [<i>tag value ...</i>]
<i>set_append</i>	:	set_append <i>set-name value</i> [<i>tag value ...</i>]
<i>set_prepend</i>	:	set_prepend <i>set-name value</i> [<i>tag value ...</i>]
<i>set_remove</i>	:	set_remove <i>set-name value</i> [<i>tag value ...</i>]
<i>setup_script</i>	:	setup_script <i>script-name</i>
<i>tag</i>	:	tag <i>tag-name</i> [<i>tag ...</i>]

```

use : use package-name [ version-tag [ access-path ] ]
      [ use-option ]

version : version version-tag

version-tag : key version-number [ key release-number [ key patch-number ] ]

use_option : -no_auto_imports
              | -auto_imports

key : letter ...

version_strategy : version_strategy version-strategy-name

version-strategy-name : best_fit
                        | best_fit_no_check
                        | first_choice
                        | last_choice
                        | keep_all

```

17.6 - The internal mechanism of cmt cvs operations

Generally, CVS does not handle queries upon the repository (such as knowing all installed modules, all tags of the modules etc..). Various tools such as CVSWeb, LXR etc. provide very powerful answers to this question, but all through a Web browser.

CMT provides a hook that can be installed within a CVS repository, based on a helper script that will be activated upon a particular CVS command, and that is able to perform some level of scan within this repository and return filtered information.

More precisely, this helper script (found in `${CMTROOT}/mgr/cmt_buildcvsinfos2.sh`) is meant to be declared within the logininfo management file (see the CVS manual for more details) as one entry named

.cmtcvsinfos, able to launch the helper script. This installation can be operated at once using the following sequence:

```
sh> cd ${CMTROOT}/mgr sh> gmake installcvs
```

This mechanism is thus fully compatible with standard remote access to the repository.

Once the helper script is installed, the mechanism operates as follows (this actually describes the algorithms installed in the CvsImplementation::show_cvs_infos method available in cmt_cvs.cxx and is transparently run when one uses the cmt cvsxxx commands):

- 1. Find a location for working with temporary files.**
This is generally deduced from the `${TMPDIR}` environment variable or in `/tmp` (or in the current directory if none of these methods apply).
- 2. There, install a directory named**
`cmtcvs/<unique-name>/cmtcvsinfos`
- 3. Then, from this directory, try to run a fake import command built as follows:**
`cvs -Q import -m cmt .cmtcvsinfos/<package-name> CMT v1`

Obviously this command is fake, since no file exist in the temporary directory we have just created.

However,

- 4. This action actually triggers the**
`cmt_buildcvsinfos2.sh` script, which simply receives in its argument the module name onto which we need information. This information is obtained by scanning the files into the repository, and an answer is built with the following syntax: `[error=error-text]`
`(1) tags=tag ... (2) branches=branch ... (3)`
`subpackages=sub-package ... (4)`

