

Cfengine Reference Manual

Auto generated, self-healing knowledge
for version 3.0.0b5

1 Cfengine 3.0.0b5 Basics

Cfengine is a suite of programs for integrated autonomic management of either individual or networked computers. It has existed as a software suite since 1993 and is published under the GNU Public Licence. This document represents cfengine version 3.0.0 of cfengine, which is a radical departure from earlier versions.

Cfengine 3 has been changed to be both a more powerful tool and a much simpler tool. Cfengine 3 is not backwards compatible with the cfengine 2 configuration language, but it interoperates with cfengine 2 so that it is "run-time compatible". This means that you can change over to version 3 slowly, with low risk and at your own speed.

With cfengine 3 you can install, configure and maintain computers using powerful hands-free tools. You can also integrate knowledge management and diagnosis into the processes.

Cfengine differs from most management systems in being

- Open software (GPL).
- Lightweight and generic.
- Non-reliant on a working network to function correctly.
- Capable of making each and every host autonomous

Cfengine 3 consists of a number of components:

<code>cf-agent</code>	Active agent
<code>cf-execd</code>	Scheduler
<code>cf-graph</code>	Graph data extractor
<code>cf-know</code>	Knowledge modelling agent
<code>cf-monitord</code>	Passive monitoring agent
<code>cf-promises</code>	Promise validator
<code>cf-runagent</code>	Remote run agent
<code>cf-serverd</code>	Server agent
<code>cf-report</code>	Self-knowledge extractor

The starred components are new. The daemon formally called `cfenvd` in previous versions of cfengine is now called `cf-monitord`.

Unlike previous versions of cfengine, which had no consistent model for its features, you can recognize *everything* in cfengine 3 from just a few concepts.

Promise A statement about the state we desire to maintain

Promise bundles
 A collection of promises

Promise bodies

A part of a promise which details and constrains its nature

Data types

An interpretation of a scalar value: string, integer or real number

Variables

An association of the form "LVALUE *represents* RVALUE", where rval may be a scalar value or a list of scalar values

Functions

Built-in parameterized rvalues

Classes

Cfengine's boolean classifiers that describe context

If you have used cfengine before then the most visible part of cfengine 3 will be its new language interface. Although it has been clear for a long time that the organically grown language used in cfengines 1 and 2 developed many problems, it was not immediately clear exactly what would be better. It has taken years of research to simplify the successful features of cfengine to a single overarching model. To understand the new cfengine, it is best to set aside any preconceptions about what cfengine is today. Cfengine 3 is a genuine "next generation" effort, which is will be a springboard into the future of system management.

1.1 Why change the language?

Many attempts at improving the user interface of cfengine have been proposed but none of them have been sufficiently impressive to make the change worthwhile before now. Some have gone in for an Object Oriented approach, but this imposes a hierarchical model that does not fit cfengine's autonomous peer model. The main goal in changing the language is to simplify and improve the robustness and functionality without sacrificing the basic freedoms and concepts. Concepts such as explicit loops and and tests have long been banished from cfengine and proposals to reintroduce them have been dismissed — something better is needed. The difficulty, of course is to provide a genuine simplification and improvement that is robust and lasting: this requires a deep understanding of the problem.

Cfengine 3's new language is a direct implementation of a model developed at Oslo University College over the past four years, known colloquially as "promise theory". Promises were originally introduced by Mark Burgess as a way to talk about cfengine's model of autonomy and have since become a powerful way of modelling cooperative systems. Cfengine 3 is a generic implementation of the language of promises that allows all of the aspects of configuration management to be unified under a single umbrella.

Why talk about promises instead of simply talking about changes? After all, the trend in business and IT management today is to talk about Change Management, e.g. in the IT Infrastructure Library (ITIL) terminology. This comes from a long history of process management thinking. But we are not really interested in change – we are interested in being in a state where we don't need to make any changes. In other words we want to be able to promise that the system is correct, verify this and only make changes if our promises are not kept.

To put it another way, cfengine is not really a *change management* system, it is a *maintenance system*. Maintenance is the process of making small changes or corrections to a model. A "model" is just another word for a template or a specification of how we want the system to work. Cfengine's model is based on the idea of promises, which means that it focuses on what is stable and lasting about a system – not about what is changing.

This is an important philosophical shift. It means we are focused mainly on what is right and not on what is wrong. By saying what "right" is (the ideal state of our system) we are focussed on

the actual behaviour. If we focus too much on the changes, i.e. the differences between now and the future, we might forget to verify that what we assume is working now in fact works.

Models that talk about change management tend to forget that after every change there is a litany of *incidents* during which it is necessary to repair the system or return it to its intended state. But if we know what we have promised, it is easy to verify whether the promise is kept. This means that it is the *promises* about how the system should be that are most important, not the actual changes that are made in order to keep them.

1.2 Testing as a non-privileged user

One of the practical advantages of cfengine is that you can test it without the need for root or administrator privileges. This is recommended for all new users of cfengine 3.

Cfengine operates with the notion of a work-directory. The default work directory for the root user is `/var/cfengine` (except on Debian Linux and various derivatives which prefer `/var/lib/cfengine`). For any other user, the work directory lies in the user's home directory, named `~/cfagent`. Cfengine prefers you to keep certain files here. You should not resist this too strongly or you will make unnecessary trouble for yourself. The decision to have this 'known directory' was made to simplify a lot of configuration.

To test cfengine as an ordinary user, do the following:

- Compile and make the software.
- Copy the binaries into the work directory:

```
host$ mkdir -p ~/.cfagent/inputs
host$ mkdir -p ~/.cfagent/bin
host$ cd src
host$ cp cf-* ~/.cfagent/bin
host$ cd ../inputs
host$ cp *.cf ~/.cfagent/inputs
```

You can test the software and play with configuration files by editing the basic get-started files directly in the `~/cfagent/inputs` directory. For example, try the following:

```
host$ ~/.cfagent/bin/cf-promises
host$ ~/.cfagent/bin/cf-promises --verbose
```

This is always the way to start checking a configuration in cfengine 3. If a configuration does not pass this check/test, you will not be allowed to use it, and `cf-agent` will look for the file `failsafe.cf`.

Notice that the cfengine 3 binaries have slightly different names than the cfengine 2 binaries. They all start with the `cf-` prefix.

```
host$ ~/.cfagent/bin/cf-agent
```

1.3 The bear necessities of a cfengine 3

```
#####
#
# The starting point for every configuration
#
#####

body common control

{
any::
```

```

    bundlesequence => { "testbundle" };
}

#####

bundle agent testbundle

{
}

```

If you try to process this using the `cf-promises` command, you will see something like this:

```

atlas$ ~/LapTop/Cfengine3/trunk/src/cf-promises -f ./unit_null_config.cf
cf3:./unit_null_config.cf:21,1: syntax error, near token '}'
Summarizing promises as text to ./unit_null_config.cf.txt
Summarizing promises as html to ./unit_null_config.cf.html

```

Examine the two files produced:

```

cat ./unit_null_config.cf.txt
firefox ./unit_null_config.cf.html

```

Here is the simplest ‘Hello world’ program in cfengine 3:

```

body common control
{
bundlesequence => { "test" };
}

bundle agent test
{
reports:

Yr2008::
    "Hello world";
}

```

1.4 Familiarizing yourself

To familiarize yourself with cfengine 3, type or paste in the following example text:

```

#####
#
# Simple test execution
#
#####

body common control

{
bundlesequence => { "testbundle" };
}

#####

```

```

bundle agent testbundle

{
vars:

    "size" int => "46k";
    "rand" int => randomint("33","$(size)");

commands:

    "/bin/echo Hello"

    # additional args that do not become path of $(this.promiser)

    args => "world - $(size)/$(rand)",

    contain => standard,
    classes => mydefine("followup","alert");

followup::

    "/bin/ls"
    contain => standard;

reports:

    alert::

        "What happened?";

}

#####

body contain standard

{
exec_owner => "mark";
useshell => "true";
}

#####

body classes mydefine(class,alert)

{
on_change => { "$(class)" };
on_failure => { "$(alert)" };
}

```

If you are familiar with cfengine's history, this will look quite strange to you, but fear not.

This example shows all of the main features of cfengine: bundles, bodies, control, variables, and promises. To the casual eye it might look complex, but that is because it is explicit about all of the details. Fortunately it is easy to hide many of these details to make the example simpler without sacrificing any functionality.

The first thing to try with this example is to verify it – did we make any mistakes? Are there any inconsistencies? To do this we use the new cfengine program `cf-promises`. Let's assume that you typed this into a file called `'test.cf'` in the current directory.

```
cf-promises -f ./test.cf
```

If all is well, typing this command shows no output. Try now running the command with verbose output.

```
cf-agent -f ./test.cf -v
```

Now you see a lot of information

```
Reference time set to Sat Aug  2 11:26:06 2008

cf3 Cfengine - 3.0.0
Free Software Foundation 1994-
Donated by Mark Burgess, Oslo University College, Norway
cf3 -----
cf3 Host name is: atlas
cf3 Operating System Type is linux
cf3 Operating System Release is 2.6.22.18-0.2-default
cf3 Architecture = x86_64
cf3 Using internal soft-class linux for host linux
cf3 The time is now Sat Aug  2 11:26:06 2008
cf3 -----
cf3 Additional hard class defined as: 64_bit
cf3 Additional hard class defined as: linux_2_6_22_18_0_2_default
cf3 Additional hard class defined as: linux_x86_64
cf3 Additional hard class defined as: linux_x86_64_2_6_22_18_0_2_default
cf3 GNU autoconf class from compile time: compiled_on_linux_gnu
cf3 Interface 1: lo
cf3 Trying to locate my IPv6 address
cf3 Looking for environment from cfenvd...
cf3 Unable to detect environment from cfMonitord
-----
Loading persistent classes
-----

Loaded persistent memory
-----

cf3  > Parsing file ./test.cf
-----

Agent's basic classified context
-----

Defined Classes = ( any Saturday Hr11 Min26 Min25_30 Q2 Hr11_Q2 Day2
August Yr2008 linux atlas 64_bit linux_2_6_22_18_0_2_default x86_64
linux_x86_64 linux_x86_64_2_6_22_18_0_2_default
linux_x86_64_2_6_22_18_0_2_default__1_SMP_2008_06_09_13_53_20__0200
compiled_on_linux_gnu net_iface_lo )

Negated Classes = ( )

Installable classes = ( )
cf3 Wrote expansion summary to promise_output_common.html
cf3 Inputs are valid
```

The last two lines of this are of interest. Each time a component of cfengine 3 parses a number of promises, it summarizes the information in an HTML file called generically `promise_output_`

component-type.html. In this case the `cf-promises` command represents all possible promises, by the type "common". You can view this output file in a suitable web browser to see exactly what cfengine has understood by the configuration. The output looks something like this:

2 A simple crash course in concepts

2.1 Rules are promises

Everything in cfengine 3 can be interpreted as a promise. Promises can be made about all kinds of different subjects, from file attributes, to the execution of commands, to access control decisions and knowledge relationships.

This simple but powerful idea allows a very practical uniformity in cfengine syntax. There is only one grammatical form for statements in the language that you need to know and it looks generically like this:

```
type:
  classes::
    "promiser" -> { "promisee1", "promisee2", ... }
      attribute_1 => value_1,
      attribute_2 => value_2,
      ...
      attribute_n => value_n;
```

We speak of a promiser (the abstract object making the promise), the promisee is the abstract object to whom the promise is made, and then there is a list of associations that we call the ‘body’ of the promise, which together with the promiser-type tells us what it is all about.

Not all of these elements are necessary every time. Some promises contain a lot of implicit behaviour. In other cases we might want to be much more explicit. For example, the simplest promise looks like this:

```
commands:
  "/bin/echo hello world";
```

This promise has default attributes for everything except the ‘promiser’, i.e. the command string that promises to execute. A more complex promise contains many attributes:

```
files:
  "/home/mark/tmp/test_plain" -> "system blue team",
    comment => "This comment follows the rule for knowledge integration",
    perms   => users("@(usernames)"),
    create  => "true";
```

The list of promisees is not used by cfengine except for documentation, just as the comment attribute (which can be added to any promise) has no actual function other than to provide more information to the user in error tracing and auditing.

You see several kinds of object in this example. All literal strings (e.g. "true") in cfengine 3 must be quoted. This provides absolute consistency and makes type-checking easy and error-correction powerful. All function-like objects (e.g. users("..")) are either builtin special functions or parameterized templates which contain the ‘meat’ of the right hand side.

2.2 Containers

Cfengine allows you to group multiple promise statements into containers called bundles.

```
bundle agent identifier
{
  commands:

    "/bin/echo These commands are a silly way to use cfengine";
    "/bin/ls -l";
    "/bin/echo But they illustrate a point";
}

```

Bundles serve two purposes: they allow us to collect related promises under a single heading, like a subroutine, and they allow us to mix configuration for different parts of cfengine in the same file. The type of a bundle is the name of the component of cfengine for which it is intended.

For instance, we can make a self-contained example agent-server configuration by labelling the bundles:

```
#
# Not a complete example
#

bundle agent testbundle
{
  files:

    "/home/mark/tmp/testcopy"

    copy_from => mycopy("/home/mark/LapTop/words", "127.0.0.1"),
    perms      => system,
    depth_search => recurse("inf");
}

#

bundle server access_rules
{
  access:

    "/home/mark/LapTop"

    admit => { "127.0.0.1" };
}

```

Another type of container in cfengine 3 is a ‘body’ part. Body parts exist to hide complex parameter information in reusable containers. The right hand side of some attribute assignments use body containers to reduce the amount of in-line information and preserve readability. You cannot choose where to use bodies: either they are used or they are not used for a particular kind

of attribute. What you can choose, however, is the name and number of parameters for the body; and you can make as many of them as you like: For example:

```
body copy_from mycopy(from,server)

{
source      => "${from}";
servers     => { "${server}" };
copy_backup => "true";

special_class::

  purge     => "true";
}
```

Notice also that classes can be used in bodies as well as parameters so that you can hide environmental adaptations in these bodies also. The classes used here are effectively ANDed with the classes under which the calling promise is defined.

2.3 Types in cfengine 3

A key difference in cfengine 3 compared to earlier versions is the presence of data types. Data types are a mechanism for associating values and checking consistency in a language. Once again, there is a simple pattern to types in cfengine.

The principle is very simple: types exist in order to match like a plug-socket relationship. In the examples above, you can see two places where types are used to match templates:

- Matching bundles to components:

```
bundle TYPE name # matches TYPE to running agent
{
}
```

- Match bodies templates to lvalues in lvalues => rvalue constraints:

```
body TYPE name # matches TYPE => name in promise
{
}
```

Check these by identifying the words ‘agent’ and ‘copy_from’ in the examples above. Types are there to make configuration more robust.

3 How to run cfengine 3 examples

These instructions assume that you have all of your configuration in a single test file, such as the example in the distribution directory 'tests/units'.

1. Test the file as a non-privileged user first, if you can.
2. Always verify syntax first with `cf-promises`. This requires no privileges. An `cf-agent` will not execute a configuration that has not passed this test.

```
host$ cf-promises -f ./inputfile.cf
```

3. Run the examples like this, e.g.

```
host$ src/cf-promises -f ./tests/units/unit_server_copy_localhost.cf
host$ src/cf-serverd -f ./tests/units/unit_server_copy_localhost.cf
host$ src/cf-agent -f ./tests/units/unit_server_copy_localhost.cf
```

Running `cf-agent` in verbose mode provides detailed information about the state of the systems promises.

```
Outcome of version 1.2.3: Promises observed to be kept 99%,
Promises repaired 1%, Promises not repaired 0%
```

The logfile 'WORKDIR/promise.log' contains the summary of these reports with timestamps. This is the simplest kind of high level audit record of the system.

4 A complete configuration

To illustrate a complete configuration for agents and daemons, consider the following example. Note how an existing user of cfengine 2 can integrate this with cfengine 2.

```
#####
#
# The starting point for every configuration
#
#####

body common control

{
any::

    bundlesequence => { "update", "cfengine2", "main" };

    inputs => { "update.cf", "templates.cf" };
}

#####

body executor control

{
splaytime => "1";
mailto => "cfengine@example.org";
}

body agent control
{
domain => "example.org";
}

#####

bundle agent main

{
processes:

    "cfenvd"    process_count    => running("cfenvd");
    "cfservd"   process_count    => running("cfservd");

commands:

    "/usr/sbin/ntpdate ntpserver.domain.tld" contain => silent;
```

```

start_cfserver::

    "/var/cfengine/bin/cfservd";

start_cfenvd::

    "/var/cfengine/bin/cfenvd -H -T";

files:

"/cfengine"  perms  => sys_public;

"/etc/passwd" perms  => sys_public,
               changes => tripwire;

"/var/cfengine/outputs"
    delete => tidy,
    file_select => age("3"),
    depth_search => recurse("inf");

"/home/backup/remotehost"

    copy_from => remotehost,
    depth_search => infin,
    action => longjob;

}

#####

bundle agent cfengine2

{
  commands:

    "/var/cfengine/bin/cfagent";

}

#####

body copy_from remotehost

{
  source      => "/iu/remotehost";
}

```

```

servers      => { "remotehost.example.org" };
copy_backup => "false";
purge        => "false";
type_check   => "true";
}

#####

body depth_search infin

{
depth => "inf";
}

body action longjob
{
ifelapsed => "480";
expireafter => "480";
background => "true";
}

```

The template file looks like this:

```

#####
#
# Our re-usable library / bag of tricks
#
#####

body action parallelize
{
background => "false";
}

#####

body perms synapses
{
mode => "0644";
owners => { "mark" };
}

#####

body perms sys_public
{
mode => "0644";
}

```

```
owners => { "root" };
}
```

```
#####
```

```
body process_count running(s)
```

```
{
match_range => "1,10"; # or irange("1","10");
out_of_range_define => { "$(s)_out_of_control" };
}
```

```
#####
```

```
body changes tripwire
```

```
{
hash          => "best";
report_changes => "content";
update        => "yes";
}
```

```
#####
```

```
body contain silent
```

```
{
no_output => "true";
}
```

```
#####
```

```
body depth_search recurse(d)
```

```
{
depth => "$(d)";
}
```

```
#####
```

```
body delete tidy
```

```
{
dirlinks => "delete";
rmdirs   => "false";
}
```

```
#####
```

```
body file_select age(d)

{
mtime      => irange(ago(0,0,$(d),0,0,0),now);
file_result => "mtime";
}
```

4.1 What should a failsafe file contain?

The ‘failsafe.cf’ file is to make sure that your system can upgrade gracefully to new versions even when mistakes are made.

As a general rule:

- Upgrade the software first, then add new features to the configuration.
- Never use advanced features in the failsafe file.

A cfengine configuration will fail-over to the failsafe configuration if it is unable to read or parse the contents successfully. That means that any new features you try in a configuration will cause a fail-over, because the parser will not be able to interpret the new features until the software itself has been updated.

4.2 Recovery from errors in the configuration

The ‘failsafe.cf’ file should be able to download the latest master configuration from source always.

```
#####
#
# failsafe.cf
#
#####

body common control

{
bundlesequence => { "update" };
}

#####

bundle agent update
{
files:

    "/var/cfengine/inputs"

    perms => system,
```

```

    copy_from => mycopy("/home/mark/cfengine-inputs","localhost"),
    depth_search => recurse("inf");

"/var/cfengine/bin"

    perms => system,
    copy_from => mycopy("/usr/local/sbin","localhost"),
    depth_search => recurse("inf");

}

#####

body perms system

{
mode => "0700";
}

#####

body depth_search recurse(d)

{
depth => "$(d)";
}

#####

body file_select cf3_files

{
leaf_name => { "cf-.*" };

file_result => "leaf_name";
}

#####

body copy_from mycopy(from,server)

{
source      => "$(from)";
servers     => { "$(server)" , "failover.domain.tld" };
#backup     => "true";
#trustkey   => "true";
encrypt     => "true";
}

```

```
}

```

4.3 Recovery from errors in the software

The update should optionally include an update of software so that a single failover from a configuration that is ‘too new’ for the software will still correct itself once the new software is available.

```
#####
#
# update.cf
#
#####

bundle agent update

{
files:

  "/var/cfengine/inputs"

  perms => system("600"),
  copy_from => mycopy("/home/mark/cfengine-inputs","localhost"),
  depth_search => recurse("inf");

  "/var/cfengine/bin"

  perms => system("700"),
  copy_from => mycopy("/usr/local/sbin","localhost"),
  file_select => cf3_files,
  depth_search => recurse("inf");

}

#####

body perms system(p)

{
mode => "${p}";
}

#####

body file_select cf3_files

```

```
{
leaf_name => { "cf-.*" };

file_result => "leaf_name";
}
```

```
#####
```

```
body copy_from mycopy(from,server)
```

```
{
source      => "${from}";
compare     => "digest";
}
```


5 Control promises

While promises to configure your system are entirely user-defined, the details of the operational behaviour of the cfengine software is of course hard-coded. You can still configure the details of this behaviour using the control promise bodies. Control behaviour is defined in bodies because the actual promises are fixed and you only change their details within sensible limits.

Note that in cfengine's previous versions, the `control` part of the configuration contained a mixture of internal control parameters and user definitions. There is now a cleaner separation in cfengine. User defined behaviour requires a promise, and must therefore be defined in a bundle.

Below is a list of the control parameters for the different components (agents and daemons) of the cfengine software.

5.1 common control promises

```
body common control
{
inputs => {
    "update.cf",
    "library.cf"
};

bundlesequence => {
    update("policy_host.domain.tld"),
    "main",
    "cfengine2"
};

output_prefix => "cfengine>";
version => "1.2.3";
}
```

the `common control` body refers to those promises that are hard-coded into all the components of cfengine, and therefore affect the behaviour of all the components.

5.1.1 bundlesequence

Type: `slist`

Allowed input range: `.*`

Synopsis: List of promise bundles to verify in order

Example:

```
body common control

{
#..
bundlesequence => {
    update("policy_host.domain.tld"),
    "main",
    "cfengine2"
};
}
```

Notes:

The `bundlesequence` determines whether compiled bundles will be executed and in what order they will be executed. The list refers to the names of bundles which might be parameterized function-like objects.

5.1.2 inputs

Type: slist

Allowed input range: .*

Synopsis: List of filenames to parse for promises

Example:

```
body common control

{
inputs => {
    "update.cf",
    "library.cf"
};
}
```

Notes:**5.1.3 version**

Type: string

Allowed input range: (arbitrary string)

Synopsis: Scalar version string for this configuration

Example:

```
body common control
{
  version => "1.2.3";
}
```

Notes:

The version string is used in error messages and reports.

5.1.4 lastseenexpireafter

Type: int

Allowed input range: 0,999999999999

Synopsis: Number of minutes after which last-seen entries are purged

Example:

Notes:

```
body common control
{
  lastseenexpireafter => "72";
}
```

5.1.5 output_prefix

Type: string

Allowed input range: (arbitrary string)

Synopsis: The string prefix for standard output

Example:

```
body common control
{
  output_prefix => "my_cf3";
}
```

Notes:**5.1.6 domain**

Type: string

Allowed input range: .*

Synopsis: Specify the domain name for this host

Example:

```
body common control
{
domain => "example.org";
}
```

Notes:

There is no standard, universal or reliable way of determining the DNS domain name of a host, so it can be set explicitly to simplify discovery and name-lookup.

5.2 agent control promises

```
body agent control
{
123_456_789::

    domain => "mydomain.com";

123_456_789_111::

    auditing => "true";

any::

    fullencryption => "true";
}
```

Settings describing the details of the fixed behavioural promises made by `cf-agent`. For example:

5.2.1 maxconnections

Type: int

Allowed input range: 0,999999999999

Synopsis: Maximum number of outgoing connections to cf-serverd

Example:

```
body agent control
{
maxconnections => "1000";
}
```

Notes:

Watch out for kernel limitations for maximum numbers of open file descriptors which can limit this.

5.2.2 abortclasses

Type: slist

Allowed input range: .*

Synopsis: A list of classes which if defined lead to termination of cf-agent

Example:

```
body agent control

{
abortclasses => { "danger", "should_not_continue" };
}
```

Notes:

A list of classes that `cf-agent` will watch out for. If any of these classes becomes defined, it will cause the current execution of `cf-agent` to be aborted. This may be used for validation, for example.

5.2.3 abortbundleclasses

Type: slist

Allowed input range: .*

Synopsis: A list of classes which if defined lead to termination of current bundle

Example:

This example shows how to use the feature to validate input to a method bundle.

```
body common control

{
```

```
bundlesequence => { "testbundle" };
version => "1.2.3";
}
```

```
#####
```

```
body agent control
```

```
{
abortbundleclasses => { "invalid" };
}
```

```
#####
```

```
bundle agent testbundle
```

```
{
vars:

    "userlist" slist => { "xyz", "mark", "jeang", "jonhenrik", "thomas", "eben" };

methods:
```

```
methods:
```

```
    "any" usebundle => subtest("${userlist}");
}
```

```
#####
```

```
bundle agent subtest(user)
```

```
{
classes:

    "invalid" not => regcmp("[a-z][a-z][a-z][a-z]", "${user}");

reports:
```

```
reports:
```

```
    !invalid::

        "User name ${user} is valid at 4 letters";
}
```

Notes:

A list of classes that `cf-agent` will watch out for. If any of these classes becomes defined, it will cause the current bundle to be aborted. This may be used for validation, for example.

5.2.4 `addclasses`

Type: `slist`

Allowed input range: `.*`

Synopsis: A list of classes to be defined always in the current context

Example:

`Add classes` adds global, literal classes. The only predicates available during the control section are `hard-classes`.

`any::`

```
addclasses => { "My_Organization" }
```

`solaris::`

```
addclasses => { "some_solaris_alive", "running_on_sunshine" };
```

Notes:

Another place to make global aliases for system `hardclasses`. Classes here are added unequivocally to the system. If classes are used to predicate definition, then they must be defined in terms of global `hard classes`.

5.2.5 `agentaccess`

Type: `slist`

Allowed input range: `.*`

Synopsis: A list of user names allowed to execute `cf-agent`

Example:

```
agentaccess => { "mark", "root", "sudo" };
```

Notes:

A list of user names that will be allowed to attempt execution of the current configuration. This is mainly a sanity check rather than a security measure.

5.2.6 agentfacility

Type: (menu option)

Allowed input range:

```
LOG_USER
LOG_DAEMON
LOG_LOCAL0
LOG_LOCAL1
LOG_LOCAL2
LOG_LOCAL3
LOG_LOCAL4
LOG_LOCAL5
LOG_LOCAL6
LOG_LOCAL7
```

Synopsis: The syslog facility for cf-agent

Example:

```
agentfacility => "LOG_USER";
```

Notes:

Sets the agent's syslog facility level. See the manual pages for syslog.

5.2.7 auditing

Type: (menu option)

Allowed input range:

```
true
false
yes
no
on
off
```

Synopsis: true/false flag to activate the cf-agent audit log

Example:

```
body agent control
{
auditing => "true";
}
```


Notes:

If this is set, cfengine will perform auditing on promises in the current configuration. This means that all details surrounding the verification of the current promise will be recorded in the audit database. The database may be inspected with `cf-report`, or `cfshow` in cfengine 2.

5.2.8 binarypaddingchar

Type: string

Allowed input range: (arbitrary string)

Synopsis: Character used to pad unequal replacements in binary editing

Example:

```
body agent control
{
binarypaddingchar => "${const.0}";
}
```

Notes:

When editing binary files, it can be dangerous to replace a text string with one that is longer or shorter as byte references and jumps would be destroyed. Cfengine will therefore not allow replacements that are larger in size than the original, but shorter strings can be padded out to the same length.

5.2.9 bindtointerface

Type: string

Allowed input range: .*

Synopsis: Use this interface for outgoing connections

Example:

```
bindtointerface => "192.168.1.1";
```

Notes:

On multi-homed hosts, the server and client can bind to a specific interface for server traffic. The IP address of the interface must be given as the argument, not the device name.

5.2.10 hashpurge

Type: (menu option)

Allowed input range:

```
    true
    false
    yes
    no
    on
    off
```

Synopsis: true/false whether stored hashes for non-existent files should be removed

Example:

```
body agent control
{
hashpurge => "true";
}
```

Notes:

Looking for files that are not present relative to the change management database also provides warnings about file deletions.

5.2.11 hashupdates

Type: (menu option)

Allowed input range:

```
    true
    false
    yes
    no
    on
    off
```

Synopsis: true/false whether stored hashes are updated when change is detected in source

Example:

```
body agent control
{
hashupdates => "true";
}
```

Notes:

If 'true' the stored reference value is updated as soon as a warning message has been given. As most changes are benign (package updates etc) this is a common setting.

5.2.12 childlibpath

Type: string

Allowed input range: .*

Synopsis: LD_LIBRARY_PATH for child processes

Example:

```
body agent control
{
childlibpath => "/usr/local/lib:/usr/local/gnu/lib";
}
```

Notes:

This string may be used to set the internal LD_LIBRARY_PATH environment of the agent.

5.2.13 defaultcopytype

Type: (menu option)

Allowed input range:

```
mtime
atime
ctime
digest
hash
binary
```

Synopsis: (null)

Example:

```
body agent control
{
#...
defaultcopytype => "digest";
}
```

Notes:

Sets the global default policy for comparing source and image in copy transactions.

5.2.14 dryrun

Type: (menu option)

Allowed input range:

```
true
false
yes
no
on
off
```

Synopsis: All talk and no action mode

Example:

```
body agent control
{
dryrun => "true";
}
```

Notes:

If set in the configuration, cfengine makes no changes to a system, only reports what it needs to do.

5.2.15 editbinaryfilesize

Type: int

Allowed input range: 0,999999999999

Synopsis: Integer limit on maximum binary file size to be edited

Example:

```
body agent control
{
edibinaryfilesize => "10M";
}
```

Notes:

The limit on editing binary files should generally be higher than for text files. Note the use of units allowed in the integer type.

5.2.16 editfilesize

Type: int

Allowed input range: 0,99999999999

Synopsis: Integer limit on maximum text file size to be edited

Example:

```
body agent control
{
editfilesize => "120k";
}
```

Notes:

The global setting for the file-editing safety-net. Note the use of special units is allowed.

5.2.17 exclamation

Type: (menu option)

Allowed input range:

```
    true
    false
    yes
    no
    on
    off
```

Synopsis: true/false print exclamation marks during security warnings

Example:

```
body agent control
{
exclamation => "false";
}
```

Notes:

This affects only the output format of warnings.

5.2.18 expireafter

Type: int

Allowed input range: 0,9999999999

Synopsis: Global default for time before on-going promise repairs are interrupted

Example:

```
body action example
{
ifelapsed => "120";
expireafter => "240";
}
```

Notes:

The locking time after which cfengine will attempt to kill and restart its attempt to keep a promise.

5.2.19 files_single_copy

Type: slist

Allowed input range: (arbitrary string)

Synopsis: List of filenames to be watched for multiple-source conflicts

Example:

```
body agent control
{
single_copy => { "/etc/*.*", "/special/file" };
}
```

Notes:

This list of regular expressions will ensure that files matching the patterns of the list are never copied from more than one source during a single run of **cf-agent**. This may be considered a protection against accidental overlap of copies from diverse remote sources, or as a first-come-first-served disambiguation tool for lazy-evaluation of overlapping file-copy promises.

5.2.20 files_auto_define

Type: slist

Allowed input range: (arbitrary string)

Synopsis: List of filenames to define classes if copied

Example:

```
body agent control
{
files_auto_define => { "/etc/syslog\.c.*", "/etc/passwd" };
}
```

Notes:

Classes are automatically defined by the files that are copied. The file is named according to the prefixed ‘canonization’ of the file name. Canonization means that non-identifier characters are converted into underscores. Thus ‘/etc/passwd’ would canonize to ‘_etc_passwd’. The prefix ‘auto_’ is added to clarify the origin of the class. Thus in the example the copying of ‘/etc/passwd’ would lead to the class ‘auto__etc_passwd’ being defined automatically.

5.2.21 fullencrytion

Type: (menu option)

Allowed input range:

```
    true
    false
    yes
    no
    on
    off
```

Synopsis: Full encryption mode in server connections, includes directory listings

Example:

```
body agent control
{
fullencrytion => "true";
}
```

Notes:

If true this encrypts all levels of the queries to the server during file transfers. The default is to not encrypt all aspects, since this can slow down transfer and basically only contributes to global warming for most users.

5.2.22 hostnamekeys

Type: (menu option)

Allowed input range:

```

true
false
yes
no
on
off

```

Synopsis: true/false label ppkeys by hostname not IP address

Example:

```

body agent control
{
hostnamekeys => "true";
}

```

Notes:

Client side choice to base key associations on host names rather than IP address. This is useful for hosts with dynamic addresses.

5.2.23 ifelapsed

Type: int

Allowed input range: 0,99999999999

Synopsis: Global default for time that must elapse before promise will be rechecked

Example:

```

#local

body action example
{
ifelapsed   => "120";
expireafter => "240";
}

```



```
# global

body agent control
{
ifelapsed => "120";
}
```

Notes:

This overrides the global settings. Promises which take a long time to verify should usually be protected with a long value for this parameter. This serves as a resource ‘spam’ protection. A cfengine check could easily run every 5 minutes provided resource intensive operations are not performed on every run. Using time classes like Hr12 etc., is one part of this strategy; using `ifelapsed` is another which is not tied to a specific time.

5.2.24 inform**Type:** (menu option)**Allowed input range:**

```
    true
    false
    yes
    no
    on
    off
```

Synopsis: true/false set inform level default**Example:**

```
body agent control
{
inform => "true";
}
```

Notes:

Equivalent to or overrides the command line option ‘-I’. Sets the default output level ‘permanently’ within the class context indicated.

5.2.25 lastseen**Type:** (menu option)

Allowed input range:

```

    true
    false
    yes
    no
    on
    off

```

Synopsis: true/false record last observed time for all client-server connections (true)

Example:

reports:

```

    "Comment"

    lastseen => "10";

```

Notes:

After this time (hours) has passed, references to the external peer will be purged from this host's database.

5.2.26 lastseenexpireafter

Type: int

Allowed input range: 0,9999999999

Synopsis: Time in days after which non-responsive last-seen hosts are purged

Example:

Notes:

```

body common control
{
lastseenexpireafter => "72";
}

```

5.2.27 mountfilesystems

Type: (menu option)

Allowed input range:

```

true
false
yes
no
on
off

```

Synopsis: true/false mount any filesystems promised

Example:

```

body agent control
{
mountfilesystems => "true";
}

```

Notes:

Issues the generic command to mount file systems defined in the file system table.

5.2.28 nonalphanumfiles

Type: (menu option)

Allowed input range:

```

true
false
yes
no
on
off

```

Synopsis: true/false warn about filenames with no alphanumeric content

Example:

```

body agent control
{
nonalphanumericfiles => "true";
}

```

Notes:

This test is applied in all recursive/depth searches.

5.2.29 repchar**Type:** string**Allowed input range:** .**Synopsis:** The character used to canonize pathnames in the file repository**Example:**

```
body agent control
{
repchar => "_";
}
```

Notes:**5.2.30 default_repository****Type:** string**Allowed input range:** [/\].***Synopsis:** Path to the default file repository**Example:**

```
body agent control
{
default_respository => "/var/cfengine/repository";
}
```

Notes:

If defined the default repository is the location where versions of files altered by cfengine are stored. This should be understood in relation to the policy for ‘backup’ in copying, editing etc. If the backups are time-stamped, this becomes effective a version control repository.

5.2.31 secureinput**Type:** (menu option)**Allowed input range:**

```
true
false
yes
no
```

```
on
off
```

Synopsis: true/false check whether input files are writable by unauthorized users

Example:

```
body agent control
{
secureinput => "true";
}
```

Notes:

If this is set, the agent will not accept an input file that is not owned by a privileged user.

5.2.32 sensiblecount

Type: int

Allowed input range: 0,99999999999

Synopsis: Minimum number of files a mounted filesystem is expected to have

Example:

```
body agent control
{
sensiblecount => "20";
}
```

Notes:

5.2.33 sensiblesize

Type: int

Allowed input range: 0,99999999999

Synopsis: Minimum number of bytes a mounted filesystem is expected to have

Example:

```
body agent control
{
sensiblesize => "20K";
}
```

```
}
```

Notes:

5.2.34 skipidentify

Type: int

Allowed input range: (arbitrary string)

Synopsis: Do not send IP/name during server connection because address resolution is broken

Example:

```
body control agent
{
skipidentify => "true";
}
```

Notes:

Hosts that are not registered in DNS cannot supply reasonable credentials for a secondary confirmation of their identity to a cfengine server. This causes the agent to ignore its missing DNS credentials.

5.2.35 suspiciousnames

Type: slist

Allowed input range: List of names to warn about if found during any file search

Synopsis: (null)

Example:

```
body agent control
{
suspiciousnames => { ".mo", "lrk3", "rootkit" };
}
```

Notes:

If cfengine sees these names during recursive (depth) file searches it will warn about them.

5.2.36 syslog

Type: (menu option)

Allowed input range:

```
true
false
yes
no
on
off
```

Synopsis: true/false switches on output to syslog at the inform level

Example:

```
body agent example
{
  syslog => "true";
}
```

Notes:

5.2.37 timezone

Type: slist

Allowed input range: (arbitrary string)

Synopsis: List of allowed timezones this machine must comply with

Example:

```
body agent control
{
  timezone => { "MET", "CET", "GMT+1" };
}
```

Notes:

5.2.38 default_timeout

Type: int

Allowed input range: 0,9999999999

Synopsis: Maximum time a network connection should attempt to connect

Example:

```
body agent control
{
default_timeout => "10";
}
```

Notes:

The time is in seconds. It is not a guaranteed number, since it depends on system behaviour. under Linux, the kernel version plays a role, since not all system calls seem to respect the signals.

5.2.39 verbose

Type: (menu option)

Allowed input range:

```
    true
    false
    yes
    no
    on
    off
```

Synopsis: true/false switches on verbose standard output

Example:

```
control agent control
{
verbose => "true";
}
```

Notes:**5.3 server control promises**


```
body server control

{
allowconnects      => { "127.0.0.1" , ":::1" , ".*.example.org" };
allowallconnects  => { "127.0.0.1" , ":::1" , ".*.example.org" };

# Uncomment me under controlled circumstances
#trustkeysfrom    => { "127.0.0.1" , ":::1" , ".*.example.org" };
}
```

Server controls are mainly about determining access policy for the connection protocol: i.e. access to the server itself. Access to specific files must be granted in addition.

5.3.1 cfruncommand

Type: string

Allowed input range: [/\].*

Synopsis: Path to the cf-agent command or cf-execd wrapper for remote execution

Example:

```
body server control

{
#..
cfruncommand => "/var/cfengine/bin/cf-agent";
}
```

Notes:

It is normal for this to point to the location of `cf-agent` but it could also point to the `cf-execd`, or even another program at your own risk.

5.3.2 maxconnections

Type: int

Allowed input range: 0,99999999999

Synopsis: Maximum number of connections that will be accepted by cf-serverd

Example:

```
body agent control
{
```

```
maxconnections => "1000";
}
```

Notes:

Watch out for kernel limitations for maximum numbers of open file descriptors which can limit this.

5.3.3 denybadclocks

Type: (menu option)

Allowed input range:

```
    true
    false
    yes
    no
    on
    off
```

Synopsis: true/false accept connections from hosts with clocks that are out of sync

Example:

```
body server body
{
#..
denybadclocks => "true";
}
```

Notes:

A possible form of attack on the fileserver is to request files based on time by setting the clocks incorrectly. This option prevents connections from clients whose clocks are drifting too far from the server clock. This serves as a warning about clock asynchronization and also a protection against Denial of Service attempts based on clock corruption.

5.3.4 allowconnects

Type: slist

Allowed input range: (arbitrary string)

Synopsis: List of IPs or hostnames that may connect to the server port

Example:

```
allowconnects => {
    "127.0.0.1" ,
    "::1",
    "2001.10.*" ,
    "host.domain.tld",
    "host[0-9]+\domain.com"
};
```

Notes:

If a client's identity matches an entry in this list it is granted to permission to send data to the server port. Clients who are not in this list may not connect or send data to the server.

5.3.5 denyconnects

Type: slist

Allowed input range: (arbitrary string)

Synopsis: List of IPs or hostnames that may NOT connect to the server port

Example:

```
body server control
{
denyconnects => { "badhost.domain.evil" };
}
```

Notes:

Hosts or IP addresses that are explicitly denied access. This should only be used in special circumstances. One should never grant generic access to everything and then deny special cases. Since the default server behaviour is to grant no access to anything, this list is unnecessary unless you have already granted access to some set of hosts using a generic pattern, to which you intend to make an exception.

5.3.6 allowallconnects

Type: slist

Allowed input range: (arbitrary string)

Synopsis: List of IPs or hostnames that may have more than one connection to the server port

Example:

```
allowallconnects => {
    "127.0.0.1" ,
```

```

    "::1",
    "2001.10.*" ,
    "host.domain.tld",
    "host[0-9]+\.\domain.com"
  };

```

Notes:

This list of regular expressions matches hosts that are allowed to connect an unlimited number of times up to the maximum connection limit. Without this, a host may only connect once (which is a very strong constraint, as the host must wait for the TCP FIN_WAIT to expire before reconnection can be attempted).

In cfengine 2 this corresponds to `AllowMultipleConnectionsFrom`.

5.3.7 trustkeysfrom

Type: slist

Allowed input range: (arbitrary string)

Synopsis: List of IPs or hostnames from whom we accept public keys on trust

Example:

```

body server control
{
trustkeysfrom => {"hosts.*", "192.168.*"};
}

```

Notes:

If connecting clients' public keys have not already been trusted, this allows us to say 'yes' to accepting the keys on trust. Normally this should be an empty list except in controlled circumstances.

5.3.8 allowusers

Type: slist

Allowed input range: (arbitrary string)

Synopsis: List of usernames who may execute requests from this server

Example:

```

allowusers => { "cfengine", "root" };

```

Notes:

The usernames listed in this list are those asserted as public key identities during client-server connections. These may or may not correspond to system identities on the server-side system.

5.3.9 dynamicaddresses

Type: slist

Allowed input range: (arbitrary string)

Synopsis: List of IPs or hostnames for which the IP/name binding is expected to change

Example:

```
body server control
{
dynamicaddresses => { "dhcp_.*" };
}
```

Notes:

The addresses or hostnames here are expected to have non-permanent address-name bindings, we must therefore work harder to determine whether hosts credentials are trusted by looking for existing public keys in files that do not match the current hostname or IP.

5.3.10 skipverify

Type: slist

Allowed input range: (arbitrary string)

Synopsis: List of IPs or hostnames for which we expect no DNS binding and cannot verify

Example:

```
body server control
{
skipverify => { "special_host.*", "192.168.*" };
}
```

Notes:

Server side decision to ignore requirements of DNS identity confirmation.

5.3.11 logallconnections

Type: (menu option)

Allowed input range:

```
true
false
yes
no
on
off
```

Synopsis: true/false causes the server to log all new connections to syslog

Example:

```
body server control
{
logallconnections => "true";
}
```

Notes:

If set, the server will record connection attempts in syslog.

5.3.12 logencryptedtransfers

Type: (menu option)

Allowed input range:

```
true
false
yes
no
on
off
```

Synopsis: true/false log all successful transfers required to be encrypted

Example:

```
body server control
{
logencryptedtransfers => "true";
}
```

Notes:

If true the server will log all transfers of files which the server requires to encrypted in order to grant access (see `ifencrypted`) to syslog. These files are deemed to be particularly sensitive.

5.3.13 hostnamekeys

Type: (menu option)

Allowed input range:

```
true
false
yes
no
on
off
```

Synopsis: true/false store keys using hostname lookup instead of IP addresses

Example:

```
body agent control
{
hostnamekeys => "true";
}
```

Notes:

Client side choice to base key associations on host names rather than IP address. This is useful for hosts with dynamic addresses.

5.3.14 auditing

Type: (menu option)

Allowed input range:

```
true
false
yes
no
on
off
```

Synopsis: true/false activate auditing of server connections

Example:

```
body agent control
```

```
{
auditing => "true";
}
```

Notes:

If this is set, cfengine will perform auditing on promises in the current configuration. This means that all details surrounding the verification of the current promise will be recorded in the audit database. The database may be inspected with `cf-report`, or `cfshow` in cfengine 2.

5.3.15 bindtointerface**Type:** string**Allowed input range:** (arbitrary string)**Synopsis:** IP of the interface to which the server should bind on multi-homed hosts**Example:**

```
bindtointerface => "192.168.1.1";
```

Notes:

On multi-homed hosts, the server and client can bind to a specific interface for server traffic. The IP address of the interface must be given as the argument, not the device name.

5.3.16 serverfacility**Type:** (menu option)**Allowed input range:**

```
LOG_USER
LOG_DAEMON
LOG_LOCAL0
LOG_LOCAL1
LOG_LOCAL2
LOG_LOCAL3
LOG_LOCAL4
LOG_LOCAL5
LOG_LOCAL6
LOG_LOCAL7
```

Synopsis: Menu option for syslog facility level**Example:**


```
body server control
{
serverfacility => "LOG_USER";
}
```

Notes:

See syslog notes.

5.4 monitor control promises

```
body monitor control()
{
#version => "1.2.3.4";

threshold => "0.3";
forgetrate => "0.7";
histograms => "true";
tcpdump => "false";
tcpdumpcommand => "/usr/sbin/tcpdump -i eth1 -n -t -v";

}
```

The system defaults will be sufficient for most users. This configurability potential will be a key to developing the integrated monitoring capabilities of cfengine however.

5.4.1 forgetrate

Type: real

Allowed input range: 0,1

Synopsis: Decimal fraction [0,1] weighting of new values over old in 2d-average computation

Example:

```
body monitor control
{
threshold => "0.3";
forgetrate => "0.7";
histograms => "true";
}
```

Notes:

Configurable settings for the machine-learning algorithm that tracks system behaviour. This is only for expert users. This parameter effectively determines (together with the monitoring rate) how quickly cfengine forgets its previous history.

5.4.2 monitorfacility

Type: (menu option)

Allowed input range:

```
LOG_USER
LOG_DAEMON
LOG_LOCAL0
LOG_LOCAL1
LOG_LOCAL2
LOG_LOCAL3
LOG_LOCAL4
LOG_LOCAL5
LOG_LOCAL6
LOG_LOCAL7
```

Synopsis: Menu option for syslog facility

Example:

```
body monitor control
{
monitorfacility => "LOG_USER";
}
```

Notes:

See notes for syslog.

5.4.3 histograms

Type: (menu option)

Allowed input range:

```
true
false
yes
no
on
```

off

Synopsis: true/false store signal histogram data

Example:

```
body monitor control
{
  histograms => "true";
}
```

Notes:

This is like the '-H' option to `cfenvd` in `cfengine 2`. It causes `cfengine` to learn the conformally transformed distributions of fluctuations about the mean.

5.4.4 tcpdump

Type: (menu option)

Allowed input range:

```
true
false
yes
no
on
off
```

Synopsis: true/false use tcpdump if found

Example:

```
body monitor control
{
  tcpdump => "true";
}
```

Notes:

Interface with TCP stream if possible.

5.4.5 tcpdumpcommand

Type: string

Allowed input range: [/\].*

Synopsis: Path to the tcpdump command on this system

Example:

```
body monitor control
{
tcpdumpcommand => "/usr/sbin/tcpdump -i eth1";
}
```

Notes:

If this is defined, the monitor will try to interface with the TCP stream and monitor generic package categories for anomalies.

5.5 runagent control promises

```
body runagent control
{
# default port is 5308

hosts => { "127.0.0.1:5308", "eternity.iu.hio.no:80", "slogans.iu.hio.no" };

#output_to_file => "true";
}
```

The most important parameter here is the list of hosts that the agent will poll for connections. This is easily read in from a file list, however when doing so always have a stable input source that does not depend on the network (including a database or directory service) in any way: introducing such dependencies makes configuration brittle.

5.5.1 hosts

Type: slist

Allowed input range: (arbitrary string)

Synopsis: List of host or IP addresses to attempt connection with

Example:

```
body runagent control
{
```

```
network1::
  hosts => { "host1.example.org", "host2", "host3" };

network2::
  hosts => { "host1.example.com", "host2", "host3" };
}
```

Notes:

The complete list of contactable hosts. The values may be either numerical IP addresses or DNS names, optionally suffixed by a ':' and a port number. If no port number is given, the default cfengine port 5308 is assumed.

5.5.2 port**Type:** int**Allowed input range:** 1024,99999**Synopsis:** Default port for cfengine server**Example:**

```
body copy_from example
{
  portnumber => "5308";
}
```

Notes:

The standard or registered port number is tcp/5308. Cfengine does not presently use its registered udp port with the same number, but this could change in the future.

5.5.3 force_ipv4**Type:** (menu option)**Allowed input range:**

```

true
false
yes
no
on
off
```

Synopsis: true/false force use of ipv4 in connection

Example:

```
body copy_from example
{
force_ipv4 => "true";
}
```

Notes:

IPv6 should be harmless to most users unless you have a partially or misconfigured setup.

5.5.4 trustkey

Type: (menu option)

Allowed input range:

```
    true
    false
    yes
    no
    on
    off
```

Synopsis: true/false automatically accept all keys on trust from servers

Example:

```
body copy_from example
{
trustkey => "true";
}
```

Notes:

If the server's public key has not already been trusted, this allows us to accept the key in automated key-exchange.

5.5.5 encrypt

Type: (menu option)

Allowed input range:

```
    true
```

```
false
yes
no
on
off
```

Synopsis: true/false encrypt connections with servers

Example:

```
body copy_from example
{
servers => { "remote-host.example.org" };
encrypt => "true";
}
```

Notes:

Client connections are encrypted with using a Blowfish randomly generated session key. The initial connection is encrypted using the public/private keys for the client and server hosts.

5.5.6 background_children

Type: (menu option)

Allowed input range:

```
true
false
yes
no
on
off
```

Synopsis: true/false parallelize connections to servers

Example:

```
body runagent control
{
background_children => "true";
}
```

Notes:

Causes the runagent to attempt parallelized connections to the servers.

5.5.7 max_children**Type:** int**Allowed input range:** 0,999999999999**Synopsis:** Maximum number of simultaneous connections to attempt**Example:**

```
body runagent control
{
max_children => "10";
}
```

Notes:

The maximum number of forked background processes allowed when parallelizing connections to servers.

5.5.8 output_to_file**Type:** (menu option)**Allowed input range:**

```

true
false
yes
no
on
off
```

Synopsis: true/false whether to send collected output to file(s)**Example:**

```
body runagent control
{
output_to_file => "true";
}
```

Notes:

Filenames are chosen automatically and placed in the 'WORKDIR/outputs/hostname_runagent.out'.

5.6 executor control promises

```
body executor control
{
  splaytime => "5";
  mailto    => "cfengine@example.org";
  mailfrom  => "cfengine@$(host).example.org";
  smtpserver => "localhost";
  schedule  => { "Min00_05", "Min30_35" }
}
```

These body settings determine the behaviour of `cf-execd`, including scheduling times and output capture to `WORKDIR/outputs` and relay via email. Note that the `splaytime` and `schedule` parameters are now coded here rather than (as previously) in the agent.

5.6.1 splaytime

Type: int

Allowed input range: 0,999999999999

Synopsis: Time in minutes to splay this host based on its name hash

Example:

```
body executor control
{
  splaytime => "2";
}
```

Notes:

A rough rule of thumb for scaling of small updates is set the splay time by:

$$\text{splaytime} = 1 + \text{Number of clients} / 50$$

5.6.2 mailfrom

Type: string

Allowed input range: *.*.*

Synopsis: Email-address cfengine mail appears to come from

Example:

```
body executor control
{
mailfrom => "MrCfengine@example.org";
}
```

Notes:

5.6.3 mailto

Type: string

Allowed input range: *.*

Synopsis: Email-address cfengine mail is sent to

Example:

```
body executor control
{
mailto => "cfengine_alias@example.org";
}
```

Notes:

The address to whom email is sent if an smtp host is configured.

5.6.4 smtpserver

Type: string

Allowed input range: .*

Synopsis: Name or IP of a willing smtp server for sending email

Example:

```
body executor control
{
smtpserver => "smtp.example.org";
}
```

Notes:

This should point to a standard port 25 server without encryption. If you are running secured or encrypted email then you should run a mail relay on localhost and point this to localhost.

5.6.5 mailmaxlines

Type: int

Allowed input range: 0,1000

Synopsis: Maximum number of lines of output to send by email

Example:

```
body executor control
{
mailmaxlines => "100";
}
```

Notes:

This limit prevents anomalously large outputs from clogging up a system administrator's mailbox. The output is truncated in the email report, but the complete original transcript is stored in 'WORKDIR/outputs/*' where it can be viewed on demand. A reference to the appropriate file is given.

5.6.6 schedule

Type: slist

Allowed input range: (arbitrary string)

Synopsis: The class schedule for activating cf-execd

Example:

```
body executor control
{
schedule => { "Min00_05", "Min05_20", "Min30_35", "Min45_50" };
}
```

Notes:

The list should contain classes which are visible to the `cf-execd` daemon. In principle any defined class will cause the daemon to wake up and schedule the execution of the `cf-agent`.

5.6.7 executorfacility

Type: (menu option)

Allowed input range:

```

LOG_USER
LOG_DAEMON
LOG_LOCAL0
LOG_LOCAL1
LOG_LOCAL2
LOG_LOCAL3
LOG_LOCAL4
LOG_LOCAL5
LOG_LOCAL6
LOG_LOCAL7

```

Synopsis: Menu option for syslog facility level

Example:

```

body executor control
{
  executorfacility => "LOG_USER";
}

```

Notes:

See the syslog manual pages.

5.6.8 execcommand

Type: string

Allowed input range: [/\].*

Synopsis: The path to the command executed by default (overriding builtin)

Example:

```

body executor control
{
  execcommand => "/usr/local/sbin/cf-agent";
}

```

Notes:

The default value for the `execcommand` is `'WORKDIR/bin/cf-agent'`.

5.7 knowledge control promises

```
body knowledge control
{
  build_directory => ".";

  sql_database => "my_knowledge";
  sql_owner => "db_user";
  sql_passwd => ""; # No passwd
  sql_type => "mysql";

  query_output => "html";

  style_sheet => "http://www.example.org/css/style.css";
  html_banner =>
    "
      <ul>
        <li>Item 1
        <li>Item 2
      </ul>
    ";
}
```

These parameters control the way in which knowledge data are stored and retrieved from a relational database and the output format of the queries.

5.7.1 id_prefix

Type: string

Allowed input range: .*

Synopsis: The LTM identifier prefix used to label topic maps (used for disambiguation in merging)

Example:

```
body knowledge control
{
  id_prefix => "unique_prefix";
}
```

Notes:

Use to disambiguate indentifiers for a successful merging of topic maps, especially in Linear Topic Map (LTM) format using third party tools such as Ontopia's Omnigator.

5.7.2 build_directory

Type: string

Allowed input range: .*

Synopsis: The directory in which to generate output files

Example:

```
body knowledge control

{
#..
build_directory => "/tmp/buildddir";
}
```

Notes:

The directory where all auto-generated textual output is placed by `cf-know`. This includes manual generation, ontology and topic map data.

5.7.3 sql_type

Type: (menu option)

Allowed input range:

```
mysql
postgres
```

Synopsis: Menu option for supported database type

Example:

```
body knowledge control
{
sql_type => "mysql";
}
```

Notes:

5.7.4 sql_database

Type: string

Allowed input range: (arbitrary string)

Synopsis: Name of database used for the topic map

Example:

```
body knowledge control
{
sql_database => "cfengine_knowledge_db";
}
```

Notes:

The name of an SQL database for caching knowledge.

5.7.5 sql_owner

Type: string

Allowed input range: (arbitrary string)

Synopsis: User id of sql database user

Example:

```
body knowledge control
{
sql_owner => "db_owner";
}
```

Notes:

Part of the credentials for opening the database. This depends on the type of database.

5.7.6 sql_passwd

Type: string

Allowed input range: (arbitrary string)

Synopsis: Embedded password for accessing sql database

Example:

```
body knowledge control
{
sql_passwd => "";
}
```

Notes:

Part of the credentials for connecting to the database server. This is system dependent. If the server host is localhost a password might not be required.

5.7.7 sql_server

Type: string

Allowed input range: (arbitrary string)

Synopsis: Name or IP of database server (or localhost)

Example:

```
body knowledge control
{
  sql_server => "localhost";
}
```

Notes:

The host name of IP address of the server. The default is to look on the localhost.

5.7.8 query_output

Type: (menu option)

Allowed input range:

```
        html
        text
```

Synopsis: Menu option for generated output format

Example:

```
body knowledge control
{
  query_output => "html";
}
```

Notes:

5.7.9 query_engine

Type: string

Allowed input range: (arbitrary string)

Synopsis: Name of a dynamic web-page used to accept and drive queries in a browser

Example:

```
body knowledge control
{
query_engine => "http://www.example.org/script.ext";
}
```

Notes:

When displaying topic maps in HTML format, `cf-know` will render each topic as a link to this URL with the new topic as an argument. Thus it is possible to make a dynamic web query by embedding `cfengine` in the web page as system call and passing the argument to it.

5.7.10 style_sheet

Type: string

Allowed input range: (arbitrary string)

Synopsis: Name of a style-sheet to be used in rendering html output (added to headers)

Example:

```
body knowledge control
{
style_sheet => "http://www.example.org/css/sheet.css";
}
```

Notes:

For formatting the HTML generated output of `cf-know`.

5.7.11 html_banner

Type: string

Allowed input range: (arbitrary string)

Synopsis: HTML code for a banner to be added to rendered html after the header

Example:

```
body knowledge control
{
html_banner => "<img src=\"http://www.example.org/img/banner.png\">";
}
```

Notes:

This content is cited when generating HTML output from the knowledge agent.

5.7.12 graph_output

Type: (menu option)

Allowed input range:

```
    true
    false
    yes
    no
    on
    off
```

Synopsis: true/false generate png visualization of topic map if possible (requires lib)

Example:

```
body knowledge control

{
# fix/override -g option
graph_output => "true";
}
```

Notes:

Equivalent to the use of the '-g' option for cf-know.

5.7.13 graph_directory

Type: string

Allowed input range: [/\].*

Synopsis: Path to directory where rendered .png files will be created

Example:

```
body knowledge control
{
graph_directory => "/tmp/output";
}
```

Notes:

A separate location where the potentially large number of ‘.png’ visualizations of a knowledge representation are pre-compiled. This feature only works if the necessary graphics libraries are present.

5.7.14 generate_manual

Type: (menu option)

Allowed input range:

```
    true
    false
    yes
    no
    on
    off
```

Synopsis: true/false generate texinfo manual page skeleton for this version

Example:

```
body knowledge control
{
generate_manual => "true";
}
```

Notes:

Auto-creates a manual based on the self-documented code. As the promise syntax is extended the manual self-heals. The resulting manual is generated in Texinfo format, from which all other formats can be generated.

5.7.15 manual_source_directory

Type: string

Allowed input range: [/\\].*

Synopsis: Path to directory where raw text about manual topics is found (defaults to build_directory)

Example:

```
body knowledge control
{
manual_source => "/path/cfengine_manual_commentary";
}
```

Notes:

This is used in the self-healing documentation. The directory points to a location where the Texinfo sources for per-section commentary is maintained.

5.8 reporter control promises

```
body reporter control
{
reports => { "performance", "last_seen", "monitor_history" };
build_directory => "/tmp/nerves";
report_output => "html";
}
```

Determines a list of reports to write into the build directory. The format may be in text, html or xml format. The reporter agent `cf-report` replaces both `cfshow` and `cfenvgraph`. It no longer produces output to the console.

5.8.1 reports

Type: (option list)

Allowed input range:

```
audit
performance
all_locks
active_locks
hashes
classes
last_seen
monitor_now
monitor_history
```

Synopsis: A list of reports to generate

Example:

```
body reporter control
{
reports => { "performance", "classes" };
}
```

Notes:

A list of report types that can be generated by this agent.

5.8.2 report_output

Type: (menu option)

Allowed input range:

```
html
text
xml
```

Synopsis: Menu option for generated output format

Example:

```
body reporter control
{
report_output => "html";
}
```

Notes:

Sets the output format of embedded database reports.

5.8.3 build_directory

Type: string

Allowed input range: .*

Synopsis: The directory in which to generate output files

Example:

```
body knowledge control
```

```
{
#..
build_directory => "/tmp/build_dir";
}
```

Notes:

The directory where all auto-generated textual output is placed by `cf-know`. This includes manual generation, ontology and topic map data.

5.8.4 auto_scaling**Type:** (menu option)**Allowed input range:**

```

true
false
yes
no
on
off
```

Synopsis: true/false whether to auto-scale graph output to optimize use of space**Example:**

```
body reporter control
{
auto_scaling => "true";
}
```

Notes:

Automatic scaling is the default.

5.8.5 errorBars**Type:** (menu option)**Allowed input range:**

```

true
false
yes
no
```

```

    on
    off

```

Synopsis: true/false whether to generate error bars on graph output

Example:

```

body reporter control
{
errorBars => "true";
}

```

Notes:

The default is to produce error bars. Without error bars from cfengine's machine learning data there is no way to assess the significance of an observation about the system, i.e. whether it is normal or anomalous.

5.8.6 time_stamps

Type: (menu option)

Allowed input range:

```

    true
    false
    yes
    no
    on
    off

```

Synopsis: true/false whether to generate timestamps on the output directory

Example:

```

body reporter control
{
time_stamps => "true";
}

```

Notes:

This option is only necessary with the default build directory. This can be used to keep snapshots of the system but it will result in a lot of storage be consumed. For most purposes cfengine is programmed to forget the past at a predictable rate and there is no need to override this.

5.8.7 query_engine

Type: string

Allowed input range: (arbitrary string)

Synopsis: Name of a dynamic web-page used to accept and drive queries in a browser

Example:

```
body knowledge control
{
query_engine => "http://www.example.org/script.ext";
}
```

Notes:

When displaying topic maps in HTML format, `cf-know` will render each topic as a link to this URL with the new topic as an argument. Thus it is possible to make a dynamic web query by embedding cfengine in the web page as system call and passing the argument to it.

5.8.8 style_sheet

Type: string

Allowed input range: (arbitrary string)

Synopsis: Name of a style-sheet to be used in rendering html output (added to headers)

Example:

```
body knowledge control
{
style_sheet => "http://www.example.org/css/sheet.css";
}
```

Notes:

For formatting the HTML generated output of `cf-know`.

5.8.9 html_banner

Type: string

Allowed input range: (arbitrary string)

Synopsis: HTML code for a banner to be added to rendered html after the header

Example:


```
body knowledge control
{
html_banner => "<img src=\"http://www.example.org/img/banner.png\">";
}
```

Notes:

This content is cited when generating HTML output from the knowledge agent.

6 Bundles of common

```

bundle common globals
{
  vars:

    "global_var" string => "value";

  classes:

    "global_class" expression => "value";
}

```

Common bundles may only contain the promise types that are common to all bodies. Their main function is to define cross-component global definitions. Common bundles are observed by every agent, whereas the agent specific bundle types are ignored by components other than the intended recipient.

6.1 vars promises

Whereas most promise types are specific to a particular kind of interpretation that requires a typed interpreter (the bundle type), a number of promises can be made in any kind of bundle since they are of a generic input/output nature. These are listed below.

6.1.1 string

Type: string

Allowed input range: (arbitrary string)

Synopsis: A scalar string

Example:

vars:

```

"xxx"    string => "Some literal string...";

"yyy"    string => readfile( "/home/mark/tmp/testfile" , "33" );

```

Notes:

In cfengine previously lists were represented (as in the shell) using separated scalars, e.g. like the PATH variable. This design feature turned out to be an error of judgement which has resulted in much trouble. This is no longer supported in cfengine 3. By keeping lists an independent type many limitations have been removed.

6.1.2 int

Type: int

Allowed input range: -99999999999,9999999999

Synopsis: A scalar integer

Example:

vars:

```
"scalar" int    => "16k";
```

```
"ran"   int    => randomint(4,88);
```

```
"dim_array" int => readstringarray("array_name", "/etc/passwd", "#[^\n]*", ":", 10, 4000);
```

Notes:

Integer values may use suffices 'k', 'K', 'm', 'M', etc

'k' The value multiplied by 1000.

'K' The value multiplied by 1024.

'm' The value multiplied by 1000 * 1000.

'M' The value multiplied by 1024 * 1024.

'g' The value multiplied by 1000 * 1000 * 1000.

'G' The value multiplied by 1024 * 1024 * 1024.

'%' A percentage between 1 and 100 - mainly for use in a storage context.

The value 'inf' may also be used to represent an unlimited positive value.

6.1.3 real

Type: real

Allowed input range: -9.99999E100,9.99999E100

Synopsis: A scalar real number

Example:

vars:

```
"scalar" real    => "0.5";
```

Notes:

Real numbers are not used in many places in cfengine, but they are useful for representing probabilities and performance data.

6.1.4 slist

Type: slist

Allowed input range: (arbitrary string)

Synopsis: A list of scalar strings

Example:

vars:

```
"xxx"    slist => { "literal1", "literal2" };

"yyy"    slist => {
                readstringlist(
                    "/home/mark/tmp/testlist",
                    "#[a-zA-Z0-9 ]*",
                    "[^a-zA-Z0-9]",
                    15,
                    4000
                )
            };

"zzz"    slist => { readstringlist("/home/mark/tmp/testlist2", "#[^\n]*", ",", 5, 4000) };■
```

Notes:

6.1.5 ilist

Type: ilist

Allowed input range: -99999999999,9999999999

Synopsis: A list of integers

Example:

vars:

```
"variable_id"

    ilist => { "10", "11", "12" };
```

Notes:

Integer lists are lists of strings that are expected to be treated as integers. The typing in cfengine is dynamic, so the variable types are interchangeable.

6.1.6 rlist

Type: rlist

Allowed input range: -9.99999E100,9.99999E100

Synopsis: A list of real numbers

Example:

vars:

```
"varid" rlist => { "0.1", "0.2", "0.3" };
```

Notes:

6.1.7 policy

Type: (menu option)

Allowed input range:

```
    free
    overridable
    constant
```

Synopsis: The policy for (dis)allowing redefinition of variables

Example:

vars:

```
"varid" string => "value..."
    policy => "constant";
```

Notes:

Variables can either be allowed to change their value dynamically (be redefined) or they can be constant. The use of private variable spaces in cfengine 3 makes it unlikely that variable redefinition would be necessary in cfengine 3.

6.2 classes promises

Whereas most promise types are specific to a particular kind of interpretation that requires a typed interpreter (the bundle type), a number of promises can be made in any kind of bundle since they are of a generic input/output nature. These are listed below.

6.2.1 or

Type: clist

Allowed input range: [a-zA-Z0-9_!&|.()]+

Synopsis: Combine class sources with inclusive OR

Example:

classes:

```
"compound_test"
    or => { classmatch("linux_x86_64_2_6_22.*"), "suse_10_3" };
```

Notes:

A useful construction for writing expressions that contain special functions.

6.2.2 and

Type: clist

Allowed input range: [a-zA-Z0-9_!&|.()]+

Synopsis: Combine class sources with AND

Example:

classes:

```
"compound_class" and => { classmatch("host[0-9].*"), "Monday", "Hr02" };
```

Notes:

If an expression contains a mixture of different object types that need to be ANDed together, this list form is more convenient than providing an expression.

6.2.3 xor**Type:** clist**Allowed input range:** [a-zA-Z0-9_!&|.()]+**Synopsis:** Combine class sources with XOR**Example:**

classes:

```
"another_global" xor => { "any", "linux", "solaris"};
```

Notes:**6.2.4 dist****Type:** rlist**Allowed input range:** -9.99999E100,9.99999E100**Synopsis:** Generate a probabilistic class distribution (strategy in cfengine 2)**Example:**

classes:

```
"my_dist"
```

```
dist => { "10", "20", "40", "50" };
```

Notes:

Assign one exclusive class randomly weighted on a probability distribution. This will generate the following classes:

```
my_dist      (always)
my_dist_10   (10/120 of the time)
my_dist_20   (20/120 of the time)
my_dist_40   (40/120 of the time)
my_dist_50   (50/120 of the time)
```


This was previous called a ‘strategy’ in cfengine 2.

6.2.5 expression

Type: class

Allowed input range: [a-zA-Z0-9_!&|.()]+

Synopsis: Evaluate string expression of classes in normal form

Example:

classes:

```
"class_name" expression => "solaris|(linux.specialclass)";
```

Notes:

A way of aliasing class combinations.

6.2.6 not

Type: class

Allowed input range: [a-zA-Z0-9_!&|.()]+

Synopsis: Evaluate the negation of string expression in normal form

Example:

classes:

```
"others" not => "linux|solaris";
```

Notes:

In file editing, this negates the effect of the promiser-pattern regular expression.

6.3 reports promises

Whereas most promise types are specific to a particular kind of interpretation that requires a typed interpreter (the bundle type), a number of promises can be made in any kind of bundle since they are of a generic input/output nature. These are listed below.

6.3.1 lastseen

Type: int

Allowed input range: 0,999999999999

Synopsis: Integer time threshold in hours since current peers were last seen, report absence

Example:

reports:

```
"Comment"
```

```
lastseen => "10";
```

Notes:

After this time (hours) has passed, references to the external peer will be purged from this host's database.

6.3.2 intermittency

Type: real

Allowed input range: 0,1

Synopsis: Real number threshold [0,1] of intermittency about current peers, report above

Example:

reports:

```
"Comment"
```

```
intermittency => "0.5";
```

Notes:

Report on cfengine peers in the neighbourhood watch whose observed irregularity of connection exceeds 0.5 scaled entropy units, meaning that they show an erratic pattern of connection.

6.3.3 showstate

Type: slist

Allowed input range: (arbitrary string)

Synopsis: List of services about which status reports should be reported to standard output

Example:

reports:

```
"Comment"

showstate => {"www_in", "ssh_out", "otherprocs" };
```

Notes:

The basic list of services is:

```
'users'      Users logged in
'rootprocs'  Privileged system processes
'otherprocs' Non-privileged process
'diskfree'   Free disk on / partition
'loadavg'    % kernel load utilization
'netbiosns_in'
              netbios name lookups (in)
'netbiosns_out'
              netbios name lookups (out)
'netbiosdgm_in'
              netbios name datagrams (in)
'netbiosdgm_out'
              netbios name datagrams (out)
'netbiosssn_in'
              netbios name sessions (in)
'netbiosssn_out'
              netbios name sessions (out)
'irc_in'     IRC connections (in)
'irc_out'    IRC connections (out)
'cfengine_in'
              cfengine connections (in)
'cfengine_out'
              cfengine connections (out)
'nfsd_in'    nfs connections (in)
'nfsd_out'   nfs connections (out)
'smtp_in'    smtp connections (in)
```

<code>'smtp_out'</code>	smtp connections (out)
<code>'www_in'</code>	www connections (in)
<code>'www_out'</code>	www connections (out)
<code>'ftp_in'</code>	ftp connections (in)
<code>'ftp_out'</code>	ftp connections (out)
<code>'ssh_in'</code>	ssh connections (in)
<code>'ssh_out'</code>	ssh connections (out)
<code>'wwws_in'</code>	wwws connections (in)
<code>'wwws_out'</code>	wwws connections (out)
<code>'icmp_in'</code>	ICMP packets (in)
<code>'icmp_out'</code>	ICMP packets (out)
<code>'udp_in'</code>	UDP dgrams (in)
<code>'udp_out'</code>	UDP dgrams (out)
<code>'dns_in'</code>	DNS requests (in)
<code>'dns_out'</code>	DNS requests (out)
<code>'tcpsyn_in'</code>	TCP sessions (in)
<code>'tcpsyn_out'</code>	TCP sessions (out)
<code>'tcpack_in'</code>	TCP acks (in)
<code>'tcpack_out'</code>	TCP acks (out)
<code>'tcpfin_in'</code>	TCP finish (in)
<code>'tcpfin_out'</code>	TCP finish (out)
<code>'tcpmisc_in'</code>	TCP misc (in)
<code>'tcpmisc_out'</code>	TCP misc (out)
<code>'webaccess'</code>	Websserver hits
<code>'weberrors'</code>	Websserver errors
<code>'syslog'</code>	New log entries (Syslog)

'messages' New log entries (messages)
 'temp0' CPU Temperature 0
 'temp1' CPU Temperature 1
 'temp2' CPU Temperature 2
 'temp3' CPU Temperature 3
 'cpu' %CPU utilization (all)
 'cpu0' %CPU utilization 0
 'cpu1' %CPU utilization 1
 'cpu2' %CPU utilization 2
 'cpu3' %CPU utilization 3

6.3.4 printfile (compound body)

Type: (ext body)

'number_of_lines'

Type: int

Allowed input range: 0,9999999999

Synopsis: Integer maximum number of lines to print from selected file

Example:

```
body printfile example
{
number_of_lines => "10";
}
```

Notes:

'file_to_print'

Type: string

Allowed input range: [/\].*

Synopsis: Path name to the file that is to be sent to standard output

Example:

```
body printfile example
{
file_to_print => "/etc/motd";
```

```
number_of_lines => "10";
}
```

Notes:

Include part of a file in a report.

6.4 * promises in ‘agent’

The generic ‘*’ promises refer to any promise type in `cf-agent`. The body attributes described below can be added to any promise rule in the agent. These promises are of a generic nature and address matters of a completely general nature – how cfengine behaves as it attempts to keep a promise.

files:

```
"/etc/passwd" -> "Security team"

perms => p("644"),
action => justcheck,
comment => "This was decided in internal procedures XYZ123";
```

6.4.1 action (compound body)

Type: (ext body)

‘action’ **Type:** (menu option)

Allowed input range:

```
fix
warn
nop
```

Synopsis: Whether to repair or report about non-kept promises

Example:

The following example shows a simple use of transaction control, causing the promise to be verified as a separate background process.

```
body action background

{
background => "true";
}
```

In the following example, the action includes the definition of a class based on the actions that were performed.

```
bundle edit_line MarkNRoot
{
insert_lines:

!pw_loaded::

"/etc/passwd"

insert_type => "file",
action => defineclass("pw_loaded");

delete_lines:

pw_loaded::

"mark.*|root.*" not_matching => "true";

}

#####

body action defineclass(c)
{
promise_repaired => { "$(c)" };
persist_time => "0";
}
```

Notes:

The action settings allow general transaction control to be implemented on promise verification. Action bodies place limits on how often to verify the promise and what classes to raise in the case that the promise can or cannot be kept.

‘ifelapsed’

Type: int

Allowed input range: 0,9999999999

Synopsis: Number of minutes before next allowed assessment of promise

Example:

```
#local

body action example
{
ifelapsed  => "120";
expireafter => "240";
}

# global

body agent control
{
ifelapsed  => "120";
}
```

Notes:

This overrides the global settings. Promises which take a long time to verify should usually be protected with a long value for this parameter. This serves as a resource ‘spam’ protection. A cfengine check could easily run every 5 minutes provided resource intensive operations are not performed on every run. Using time classes like Hr12 etc., is one part of this strategy; using `ifelapsed` is another which is not tied to a specific time.

‘expireafter’

Type: int

Allowed input range: 0,99999999999

Synopsis: Number of minutes before a repair action is interrupted and retried

Example:

```
body action example
{
ifelapsed  => "120";
expireafter => "240";
}
```

Notes:

The locking time after which cfengine will attempt to kill and restart its attempt to keep a promise.

'log_string'

Type: string

Allowed input range: (arbitrary string)

Synopsis: A message to be written to the log when the promise is verified

Example:

```
promise-type:

    "promiser"

    attr => "value",
    action => log_me("checked $(this.promiser)");

# ..

body action log_me(s)
{
    log_string => "$(s)";
}
```

Notes:

'log_level'

Type: (menu option)

Allowed input range:

```
inform
verbose
error
log
```

Synopsis: The reporting level sent to syslog

Example:

```
body action example
{
    log_level => "inform";
}
```

Notes:

Use this as an alternative to auditing to use the syslog mechanism to centralize or manage messaging from cfengine. A backup of these messages will still be kept in 'WORKDIR/outputs' if you are using `cf-execd`.

'audit'

Type: (menu option)**Allowed input range:**

```

true
false
yes
no
on
off

```

Synopsis: true/false switch for detailed audit records of this promise

Example:

```

body action example
{
# ...

audit => "true";
}

```

Notes:

If this is set, cfengine will perform auditing on this specific promise. This means that all details surrounding the verification of the current promise will be recorded in the audit database. The database may be inspected with `cf-report`, or `cfshow` in cfengine 2.

'background'

Type: (menu option)**Allowed input range:**

```

true
false
yes
no
on

```

off

Synopsis: true/false switch for parallelizing the promise repair

Example:

```
body action example
{
background => "true";
}
```

Notes:

If possible, background the verification of the current promise. This is advantageous only if the verification might take a significant amount of time, e.g. in remote copying of filesystem/disk scans.

'report_level'

Type: (menu option)

Allowed input range:

```
inform
verbose
error
log
```

Synopsis: The reporting level for standard output

Example:

```
body action example
{
report_level => "verbose";
}
```

Notes:

In cfengine 2 one would say 'inform=true' or 'syslog=true', etc. This replaces these levels since they act as encapsulating super-sets.

6.4.2 classes (compound body)

Type: (ext body)

'promise_repaired'**Type:** slist**Allowed input range:** [a-zA-Z0-9_\$.]+**Synopsis:** A list of classes to be defined**Example:**

```
body classes example
{
promise_repaired => { "change_happened" };
}
```

Notes:

If a promise is 'repaired' it means that a corrective action had to be taken to keep the promise.

'repair_failed'**Type:** slist**Allowed input range:** [a-zA-Z0-9_\$.]+**Synopsis:** A list of classes to be defined**Example:**

```
body classes example
{
repair_failed => { "unknown_error" };
}
```

Notes:

A promise could not be repaired because the corrective action failed for some reason.

'repair_denied'**Type:** slist**Allowed input range:** [a-zA-Z0-9_\$.]+**Synopsis:** A list of classes to be defined**Example:**

```
body classes example
{
repair_denied => { "permission_failure" };
}
```

Notes:

A promise could not be kept because access to a key resource was denied.

`'repair_timeout'`

Type: slist

Allowed input range: [a-zA-Z0-9_\$.]+

Synopsis: A list of classes to be defined

Example:

```
body classes example
{
repair_timeout => { "too_slow", "did_not_wait" };
}
```

Notes:

A promise maintenance repair timed-out waiting for some dependent resource.

`'promise_kept'`

Type: slist

Allowed input range: [a-zA-Z0-9_\$.]+

Synopsis: A list of classes to be defined

Example:

```
body classes example
{
promise_kept => { "success", "kaplah" };
}
```

Notes:

This class is set if no action was necessary by `cf-agent` because the promise concerned was already kept without further action required.

`'persist_time'`

Type: int

Allowed input range: 0,9999999999

Synopsis: A number of minutes the specified classes should remain active

Example:

```
body classes example
{
persist_time => "10";
}
```

Notes:

By default classes are ephemeral entities that disappear when `cf-agent` terminates. By setting a persistence time, they can last even when the agent is not running.

`'timer_policy'`

Type: (menu option)

Allowed input range:

```
absolute
reset
```

Synopsis: Whether a persistent class restarts its counter when rediscovered

Example:

```
body classes example
{
timer_policy => "reset";
}
```

Notes:

The in most cases resetting a timer will give a more honest appraisal of which classes are currently important, but if we want to activate a response of limited duration as a rare event then an absolute time limit is useful.

6.4.3 ifvarclass

Type: string

Allowed input range: (arbitrary string)

Synopsis: Extended classes ANDed with context

Example:

The generic example has the form:

```
promise-type:

    "promiser"

    ifvarclass => "$({program})_running|({program})_notfound&Hr12)";
```

A specific example would be:

bundle agent example

```
{
commands:

any::

    "/bin/echo This is linux"

    ifvarclass => "linux";

    "/bin/echo This is solaris"

    ifvarclass => "solaris";

}
```

Notes:

This is an additional class expression that will be evaluated after the ‘`class::`’ classes have selected promises. It is provided in order to enable a channel between variables and classes. The result is thus the logical AND of the ordinary classes and the variable classes.

This function is provided so that one can form expressions that link variables and classes, e.g.

```
# Check that all components are running
```

```
vars:
```

```

"component" slist => { "cf-monitor", "cf-serverd" };

processes:

"${component}" restart_class => canonify("start_${component}");

commands:

"/var/cfengine/bin/${component}"

    ifvarclass => canonify("start_${component}");

```

Notice that the function `canonify()` is provided to convert a general variable input into a string composed only of legal characters, using the same algorithm that cfengine uses.

6.4.4 comment

Type: string

Allowed input range: (arbitrary string)

Synopsis: Retained comment about this promise's real intention

Example:

```
comment => "This comment follows the data for reference ...",
```

Notes:

Comments written in code follow the program, they are not merely discarded. They appear in reports and error messages.

6.5 * promises in 'edit_line'

The generic '*' promises refer to any promise type in `cf-agent`. The body attributes described below can be added to any promise rule in the agent. These promises are of a generic nature and address matters of a completely general nature – how cfengine behaves as it attempts to keep a promise.

files:

```
"/etc/passwd" -> "Security team"
```



```
perms => p("644"),
action => justcheck,
comment => "This was decided in internal procedures XYZ123";
```

6.5.1 select_region (compound body)

Type: (ext body)

`'select_start'`

Type: string

Allowed input range: .*

Synopsis: Regular expression matching start of edit region

Example:

```
body select_region example(x)
```

```
{
select_start => "\[$(x)\]";
select_end => "\[.*\]";
}
```

Notes:

See also `select_end`. These delimiters mark out the region of a file to be edited. In the example, it is assumed that the file has section marks

```
[section 1]
```

```
lines.
lines...
```

```
[section 2]
```

```
lines ....
etc..
```

`'select_end'`

Type: string

Allowed input range: .*

Synopsis: Regular expression matches end of edit region from start

Example:

```
body select_region example(x)

{
select_start => "\[$(x)\]";
select_end => "\[.*\]";
}
```

Notes:

See also `select_start`. These delimiters mark out the region of a file to be edited. In the example, it is assumed that the file has section marks

```
[section 1]
```

```
lines.
lines...
```

```
[section 2]
```

```
lines ....
etc..
```

7 Bundles of agent

```

bundle agent main(parameter)
{
vars:

  "sys_files"    slist      => {
                        "/etc/passwd",
                        "/etc/services"
                      };
files:

  "${sys_files}" perms      => p("root", "0644"),
                    changes  => trip_wire;

  "/etc/shadow"  perms      => p("root", "0600"),
                    changes  => trip_wire;

  "/usr"         changes    => trip_wire,
                    depth_search => recurse("inf");

  "/tmp"         delete     => tidy,
                    file_select => days("2"),
                    depth_search => recurse("inf");
}

```

Agent bundles contain user-defined promises for `cf-agent`. The types of promises and their corresponding bodies are detailed below.

7.1 commands promises in ‘agent’

```

commands:

  "/path/to/command args"

    args => "more args",
    contain => contain_body,
    module => true/false;

```

Command *containment* allows you to make a ‘sandbox’ around a command, to run it as a non-privileged user inside an isolated directory tree. Cfengine `modules` are commands that support a simple protocol (see below) in order to set additional variables and classes on execution from user

defined code. Modules are intended for use as system probes rather than additional configuration promises.

In cfengine 3 commands and processes have been separated cleanly. Restarting of processes must be coded as a separate command. This stricter type separation will allow more careful conflict analysis to be carried out.

Output from commands executed here is quoted inline, but prefixed with the letter 'Q' to distinguish it from other output, e.g. from `reports`.

Commands were called `shellcommands` in cfengine previously. Modules were

NOTE: a common mistake in using cfengine is to embed many shell commands instead of using the built-in functionality. Use of cfengine internals is preferred as it assures convergence and proper integrated checking. Extensive use of shell commands will make a cfengine execution very heavy-weight like other management systems. To minimize the system cost of execution, always use cfengine internals.

```
bundle agent example
```

```
{
commands:

  "/bin/sleep 10"
    action => background;

  "/bin/sleep"
    args => "20",
    action => background;
}
```

7.1.1 args

Type: string

Allowed input range: (arbitrary string)

Synopsis: Alternative string of arguments for the command (concatenated with promiser string)

Example:

```
commands:

  "/bin/echo one"
```

```
args => "two three";
```

Notes:

Sometimes it is convenient to separate the arguments to a command from the command itself. The final arguments are the concatenation with one space. So in the example above the command would be

```
/bin/echo one two three
```

7.1.2 contain (compound body)

Type: (ext body)

'useshell' **Type:** (menu option)

Allowed input range:

```

true
false
yes
no
on
off
```

Synopsis: true/false embed the command in a shell environment (true)

Example:

```

body contain example
{
useshell => "true";
}
```

Notes:

The default is to use a shell when executing commands, but this has both resource and security consequences. A shell consumes an extra process and inherits environment variables, reads commands from files and performs other actions beyond the control of cfengine. If one does not need shell functionality such as piping through multiple commands then it is best to manage without it.

'umask' **Type:** (menu option)

Allowed input range:

0
77
22
27
72

Synopsis: The umask value for the child process

Example:

```
body contain example
{
umask => "077";
}
```

Notes:

Sets the internal umask for the process.

'exec_owner'

Type: string

Allowed input range: (arbitrary string)

Synopsis: The user name or id under which to run the process

Example:

```
body contain example
{
exec_owner => "mysql_user";
}
```

Notes:

This is part of the restriction of privilege for child processes when running `cf-agent` as the root user, or a user with privileges.

'exec_group'

Type: string

Allowed input range: (arbitrary string)

Synopsis: The group name or id under which to run the process

Example:

```
body contain example
{
exec_group => "nogroup";
}
```

Notes:

This is part of the restriction of privilege for child processes when running `cf-agent` as the root user, or a user with privileges.

`'exec_timeout'`

Type: int

Allowed input range: 1,3600

Synopsis: Timeout in seconds for command completion

Example:

```
body contain example
{
exec_timeout => "30";
}
```

Notes:

Attempt to time-out after this number of seconds. This cannot be guaranteed as not all commands are willing to be interrupted in case of failure.

`'chdir'`

Type: string

Allowed input range: [/\].*

Synopsis: Directory for setting current/base directory for the process

Example:

```
body contain example

{
chdir => "/containment/directory";
}
```

Notes:

This command has the effect of placing the running command into a current working directory equal to the parameter given, i.e. it works like the 'cd' shell command.

'chroot'

Type: string**Allowed input range:** [/\] .***Synopsis:** Directory of root sandbox for process**Example:**

```
body contain example
```

```
{
chroot => "/private/path";
}
```

Notes:

Sets the path of the directory that will be experienced as the top-most root directory for the process. In security parlance, this creates a 'sandbox' for the process.

'preview'

Type: (menu option)**Allowed input range:**

```
true
false
yes
no
on
off
```

Synopsis: true/false preview command when running in dry-run mode (with -n)**Example:**

```
body contain example
```

```
{
preview => "true";
}
```


Notes:

Previewing shell scripts during a dry-run is a potentially misleading activity. It should only be used on scripts that make no changes to the system. It is cfengine best practice to never write change-functionality into user-written scripts except as a last resort: cfengine can apply its safety checks to user defined scripts.

`'no_output'`

Type: (menu option)

Allowed input range:

```

true
false
yes
no
on
off

```

Synopsis: true/false discard all output from the command

Example:

```

body contain example
{
no_output => "true";
}

```

Notes:

This is equivalent to piping standard output and error to `‘/dev/null’`.

7.1.3 module

Type: (menu option)

Allowed input range:

```

true
false
yes
no
on
off

```

Synopsis: true/false whether to expect the cfengine module protocol

Example:

commands:

```
"/masterfiles/user_script"

module => "true";
```

Notes:

If true, the module protocol is supported for this script, i.e. it is treated as a user module. A plug-in module may be written in any language, it can return any output you like, but lines which begin with a '+' sign are treated as classes to be defined (like '-D'), while lines which begin with a '-' sign are treated as classes to be undefined (like '-N'). Lines starting with '=' are variables/macros to be defined. Any other lines of output are cited by cfagent, so you should normally make your module completely silent. Here is an example module written in perl.

```
#!/usr/bin/perl
#
# module:myplugin
#

# lots of computation....

if (special-condition)
{
  print "+specialclass";
}
```

Modules inherit the environment variables from cfagent and accept arguments, just as a regular command does.

```
#!/bin/sh
#
# module:myplugin
#

/bin/echo $*
```

cf-agent defines the classes as an environment variable so that programs have access to these. E.g. try the following module:

```
#!/usr/bin/perl

print "Decoding $ENV{CFALLCLASSES}\n";

@allclasses = split (":", "$ENV{CFALLCLASSES}");

while ($c=shift(@allclasses))
{
  $classes{$c} = 1;
  print "$c is set\n";
}
```

```
}
```

Modules define variables in `cf-agent` by outputting strings of the form

```
=variablename=value
```

These variables end up in a context which has the same name as the module. When the `$(allclasses)` variable becomes too large to manipulate conveniently, you can access the complete list of currently defined classes in the file `'/var/cfengine/state/allclasses'`.

```
bundle agent main(parameter)
{
vars:
  "sys_files"    slist      => {
                        "/etc/passwd",
                        "/etc/services"
                      };
files:
  "$(sys_files)" perms      => p("root","0644"),
                    changes => trip_wire;
  "/etc/shadow"  perms      => p("root","0600"),
                    changes => trip_wire;
  "/usr"         changes    => trip_wire,
                    depth_search => recurse("inf");
  "/tmp"         delete     => tidy,
                    file_select  => days("2"),
                    depth_search => recurse("inf");
}
```

Agent bundles contain user-defined promises for `cf-agent`. The types of promises and their corresponding bodies are detailed below.

7.2 files promises in 'agent'

Files promises are an umbrella concept for all attributes of files. Operations fall basically into three categories: create, delete and edit.

```
files:

  "/path/file_object"

    perms => perms_body,
    ... ;
```

Prior to version 3, file promises were scattered into many different types such as `files`, `tidy`, `copy`, `links`, etc. These are all parameterized attributes of a single type of promise now, because the promiser-object of all those matters refers to the same type of object. This greatly simplifies the matter of detecting promise conflicts.

```
bundle agent example
{
files:

  "/home/mark/tmp" -> "Security team"

    changes      => lay_a_tripwire,
    depth_search => recurse("inf"),
    action       => background;
}

#####

body changes lay_a_tripwire

{
hash          => "md5";
report_changes => "content";
update        => "yes";
}
```

7.2.1 file_select (compound body)

Type: (ext body)

'leaf_name'

Type: slist

Allowed input range: (arbitrary string)

Synopsis: List of regexes that match an acceptable name

Example:

```
body file_select example
{
leaf_name => { "S[0-9]+[a-zA-Z]+", "K[0-9]+[a-zA-Z]+" };
file_result => "leaf_name";
}
```

Notes:

This pattern matches only the node name of the file, not its path.

`'path_name'`

Type: slist

Allowed input range: [/\].*

Synopsis: List of pathnames to match acceptable target

Example:

```
body file_select example
{
leaf_name => { "prog.pid", "prog.log" };
path_name => { "/etc/*", "/var/run/*" };

file_result => "leaf_name.path_name"
}
```

Notes:

Path name and leaf name can be conveniently tested for separately by use of appropriate regular expressions.

`'search_mode'`

Type: string

Allowed input range: [0-7augorwxst,+~]+

Synopsis: Mode mask for acceptable files

Example:

```
body file_select example
{
search_mode => "644";
```

```
}

```

Notes:

The mode may be specified in symbolic or numerical form with '+' and '-' constraints.

'search_size'

Type: irange [int,int]

Allowed input range: 0,inf

Synopsis: Integer range of file sizes

Example:

```
body file_select example
{
search_size => irange("0","20k");
file_result => "size";
}
```

Notes:

'search_owners'

Type: slist

Allowed input range: [a-zA-Z0-9_\$.]+

Synopsis: List of acceptable user names or ids for the file

Example:

```
body file_select example
{
search_owners => { "mark", "jeang", "student_.*" };
file_result => "owner";
}
```

Notes:

A list of regular expressions.

'search_groups'

Type: slist

Allowed input range: [a-zA-Z0-9_\$.]+

Synopsis: List of acceptable group names or ids for the file

Example:

```
body file_select example
{
search_group => { "users", "special_.*" };
file_result => "group";
}
```

Notes:

'ctime' **Type:** irange [int,int]

Allowed input range: 0,4026531839

Synopsis: Range of change times (ctime) for acceptable files

Example:

```
body files_select example
{
ctime => irange(ago(1,0,0,0,0,0),now);
file_result => "ctime";
}
```

Notes:

The file's change time refers to both modification of content and attributes such as permissions.

'mtime' **Type:** irange [int,int]

Allowed input range: 0,4026531839

Synopsis: Range of modification times (mtime) for acceptable files

Example:

```
body files_select example
{
mtime => irange(ago(1,0,0,0,0,0),now);
}
```

```
file_result => "mtime";
}
```

Notes:

The file's modification time refers to both modification of content but not other attributes such as permissions.

'atime'

Type: irange [int,int]**Allowed input range:** 0,4026531839**Synopsis:** Range of access times (atime) for acceptable files**Example:**

```
body file_select
```

```
{
# files accessed in the last hour

atime    => irange(ago(0,0,0,1,0,0),now);
file_result => "atime";
}
```

```
body file_select
```

```
{
# files accessed since 00:00 1st Jan 2000

atime    => irange(on(2000,1,1,0,0,0),now);
file_result => "atime";
}
```

Notes:

A range of times during which a file was accessed can be specified in a `file_select` body. (Like file filters in cfengine 2.)

'exec_regex'

Type: string**Allowed input range:** [/\].***Synopsis:** Matches file if this regular expression matches any full line returned by the command

Example:

```
body file_select example
{
exec_regex => "SPECIAL_LINE: .*";
exec_program => "/path/test_program ${this.promiser}";
file_result => "exec_program.exec_regex";
}
```

Notes:

The regular expression must be used in conjunction with the `exec_program` test. In this way the program must both return exit status 0 and its output must match the regular expression.

`'exec_program'`

Type: string

Allowed input range: `[/\].*`

Synopsis: Execute this command on each file and match if the exit status is zero

Example:

```
body file_select example
{
exec_program => "/path/test_program ${this.promiser}";
file_result => "exec_program";
}
```

Notes:

This is part of the customizable file search criteria. If the user-defined program returns exit status 0, the file is considered matched.

`'filetypes'`

Type: (option list)

Allowed input range:

```
plain
reg
symlink
dir
```

```

socket
fifo
door
char
block

```

Synopsis: List of acceptable file types from menu choices

Example:

```

body file_select
{
filetypes => { "plain", "symlink" };
}

```

Notes:

'issymlinkto'

Type: slist

Allowed input range: (arbitrary string)

Synopsis: List of regular expressions to match file objects

Example:

```

body file_select example
{
issymlinkto => { "/etc/[^/]*", "/etc/init.d/[a-z0-9]*" };
}

```

Notes:

A list of regular expressions. If the file is a symbolic link which points to files matched by one of these expressions, the file will be selected.

'file_result'

Type: string

Allowed input range: [(leaf_name|path_name|file_types|mode|size|owner|group|atime|ctime|mtime|regex|exec_program)[!&!.]*

Synopsis: Logical expression combining classes defined by file search criteria

Example:

```

body file_select any_age

{
mtime      => irange(ago(1,0,0,0,0,0),now);
file_result => "mtime";
}

body file_select pdf_files_1dayold

{
mtime      => irange(ago(0,0,1,0,0,0),now);
leaf_name  => { ".*.pdf" , ".*.fdf" };

file_result => "leaf_name&mtime";
}

```

Notes:

Sets the criteria for file selection outcome during file searches. The syntax is the same as for a class expression since the file selection is a classification of the file-search in the same way that system classes are a classification of the abstract host-search.

7.2.2 copy_from (compound body)

Type: (ext body)

'source' **Type:** string

Allowed input range: [/\].*

Synopsis: Reference source file from which to copy

Example:

```

body copy_from example
{
source => "/path/to/source";
}

```

Notes:

For remote copies this refers to the file name on the remote server.

'servers' **Type:** slist

Allowed input range: (arbitrary string)

Synopsis: List of servers in order of preference from which to copy

Example:

```
body copy_from example
{
servers => { "primary.example.org", "secondary.example.org",
            "tertiary.other.domain" };
}
```

Notes:

The servers are tried in order until one of them succeeds.

'portnumber'

Type: int

Allowed input range: 1024,99999

Synopsis: Port number to connect to on server host

Example:

```
body copy_from example
{
portnumber => "5308";
}
```

Notes:

The standard or registered port number is tcp/5308. Cfengine does not presently use its registered udp port with the same number, but this could change in the future.

'copy_backup'

Type: (menu option)

Allowed input range:

```
    true
    false
    timestamp
```

Synopsis: Menu option policy for file backup/version control

Example:

```
body copy_from example
{
copy_backup => "timestamp";
}
```

Notes:

Determines whether a backup of the previous version is kept on the system. This should be viewed in connection with the system repository, since a defined repository affects the location at which the backup is stored.

'stealth' **Type:** (menu option)

Allowed input range:

```
    true
    false
    yes
    no
    on
    off
```

Synopsis: true/false whether to preserve time stamps on copied file

Example:

```
body copy_from example
{
stealth => "true";
}
```

Notes:

Preserves file access and modification times on the promiser files.

'preserve' **Type:** (menu option)

Allowed input range:

```
    true
    false
    yes
    no
    on
    off
```

Synopsis: true/false whether to preserve file permissions on copied file

Example:

```
body copy_from example
{
preserve => "true";
}
```

Notes:

Whether or not the copy preserves the permissions on the source files.

'linkcopy_patterns'

Type: slist

Allowed input range: (arbitrary string)

Synopsis: List of patterns matching symbolic links that should be replaced with copies

Example:

```
body copy_from mycopy(from)

{
source          => "$(from)";
linkcopy_patterns => { ".*" };
}
```

Notes:

The pattern matches the last node filename (i.e. without the absolute path).

'copylink_patterns'

Type: slist

Allowed input range: (arbitrary string)

Synopsis: List of patterns matching files that should be linked instead of copied

Example:

```
body copy_from example
{
```

```
copylink_patterns => { "special_node1", "other_node.*" };
}
```

Notes:

The matches are performed on the last node of the filename, i.e. the file without its path.

'compare' **Type:** (menu option)

Allowed input range:

```
    atime
    mtime
    ctime
    digest
    hash
```

Synopsis: Menu option policy for comparing source and image file attributes

Example:

```
body copy_from example

{
compare => "digest";
}
```

Notes:

The default copy method is **'mtime'** (modification time), meaning that the source file is copied to the destination (promiser) file, if the source file has been modified more recently than the destination.

'link_type'

Type: (menu option)

Allowed input range:

```
    symlink
    hardlink
    relative
    absolute
    none
```

Synopsis: Menu option for type of links to use when copying

Example:

```
body link_from example
{
link_type => "hard";
}
```

Notes:

What kind of link should be used to link files. Users are advised to be wary of ‘hard links’ (see Unix manual pages for the ‘ln’ command). The behaviour of non-symbolic links is often precarious and unpredictable.

‘type_check’

Type: (menu option)

Allowed input range:

```
    true
    false
    yes
    no
    on
    off
```

Synopsis: true/false compare file types before copying and require match

Example:

```
body copy_from example
{
type_check => "false";
}
```

Notes:

File types at source and destination should normally match in order for updates to overwrite them. This option allows this checking to be switched off.

‘force_update’

Type: (menu option)

Allowed input range:


```

true
false
yes
no
on
off

```

Synopsis: true/false force copy update always

Example:

```

body copy_from example
{
force_update => "true";
}

```

Notes:

Warning: this is a non-convergent operation. Although the end point might stabilize in content, the operation will never quiesce. Use of this feature is not recommended except in exceptional circumstances since it creates a busy-dependency. If the copy is a network copy, the system will be disturbed by network disruptions.

'force_ipv4'

Type: (menu option)

Allowed input range:

```

true
false
yes
no
on
off

```

Synopsis: true/false force use of ipv4 on ipv6 enabled network

Example:

```

body copy_from example
{
force_ipv4 => "true";
}

```

Notes:

IPv6 should be harmless to most users unless you have a partially or misconfigured setup.

'copy_size'

Type: irange [int,int]

Allowed input range: 0,inf

Synopsis: Integer range of file sizes that may be copied

Example:

```
body copy_from example
{
copy_size => irange("0","50000");
}
```

Notes:

The use of the irange function is optional. Ranges may also be specified as a comma separated numbers.

'trustkey' **Type:** (menu option)

Allowed input range:

```
true
false
yes
no
on
off
```

Synopsis: true/false trust public keys from remote server if previously unknown

Example:

```
body copy_from example
{
trustkey => "true";
}
```

Notes:

If the server's public key has not already been trusted, this allows us to accept the key in automated key-exchange.

'encrypt' **Type:** (menu option)

Allowed input range:

```

true
false
yes
no
on
off

```

Synopsis: true/false use encrypted data stream to connect to remote host

Example:

```

body copy_from example
{
servers => { "remote-host.example.org" };
encrypt => "true";
}

```

Notes:

Client connections are encrypted with using a Blowfish randomly generated session key. The initial connection is encrypted using the public/private keys for the client and server hosts.

'verify' **Type:** (menu option)

Allowed input range:

```

true
false
yes
no
on
off

```

Synopsis: true/false verify transferred file by hashing after copy (resource penalty)

Example:

```
body copy_from example
{
verify => "true";
}
```

Notes:

This is a highly resource intensive option, not recommended for large file transfers.

‘purge’

Type: (menu option)

Allowed input range:

```
    true
    false
    yes
    no
    on
    off
```

Synopsis: true/false purge files on client that do not match files on server when depth_search

Example:

```
body copy_from example
{
purge => "true";
}
```

Notes:

Purging files is a potentially dangerous matter during a file copy it implies that any promiser (destination) file which is not matched by a source will be deleted. Since there is no source, this means the file will be irretrievable. Great care should be exercised when using this feature.

‘check_root’

Type: (menu option)

Allowed input range:

```
    true
    false
    yes
```

```
no
on
off
```

Synopsis: true/false check permissions on the root directory when depth_search

Example:

```
body copy_from example
{
check_root => "true";
}
```

Notes:

When copying files recursively (by depth search), this flag determines whether the permissions of the root directory should be set from the root of the source. The default is to check only copied file objects and subdirectories within this root (false).

‘findertype’

Type: (menu option)

Allowed input range:

```
MacOSX
```

Synopsis: Menu option for default finder type on MacOSX

Example:

```
body copy_from example
{
findertype => "MacOSX";
}
```

Notes:

This applies only to the Macintosh OSX variants.

7.2.3 link_from (compound body)

Type: (ext body)

‘source’ **Type:** string

Allowed input range: (arbitrary string)

Synopsis: The source file to which the link should point

Example:

```
body copy_from example
{
source => "/path/to/source";
}
```

Notes:

For remote copies this refers to the file name on the remote server.

‘link_type’

Type: (menu option)

Allowed input range:

```
symlink
hardlink
relative
absolute
none
```

Synopsis: The type of link used to alias the file

Example:

```
body link_from example
{
link_type => "hard";
}
```

Notes:

What kind of link should be used to link files. Users are advised to be wary of ‘hard links’ (see Unix manual pages for the ‘ln’ command). The behaviour of non-symbolic links is often precarious and unpredictable.

‘copy_patterns’

Type: slist

Allowed input range: (arbitrary string)

Synopsis: A set of patterns that should be copied and synchronized instead of linked

Example:

```
body link_from example
{
copy_patterns => { "special_node1", "/path/special_node2" };
}
```

Notes:

During the linking of files, it is sometimes useful to buffer changes with an actual copy, especially if the link is to an ephemeral file system. This list of patterns matches files that arise during a linking policy. A positive match means that the file should be copied and updated by modification time.

'when_no_source'

Type: (menu option)

Allowed input range:

```
force
delete
nop
```

Synopsis: Behaviour when the source file to link to does not exist

Example:

```
body link_from example
{
when_no_file => "force";
}
```

Notes:

If we try to create a link to a file that does not exist a link, how should cfengine respond? The options are to force the creation to a file that does not (yet) exist, delete any existing link, or do nothing.

'link_children'

Type: (menu option)

Allowed input range:

```

true
false
yes
no
on
off

```

Synopsis: true/false whether to link all directory's children to source originals

Example:

```

body link_from example
{
link_children => "true";
}

```

Notes:

If the promiser is a directory, instead of copying the children, link them to the source.

'when_linking_children'

Type: (menu option)

Allowed input range:

```

override_file
if_no_such_file

```

Synopsis: Policy for overriding existing files when linking directories of children

Example:

```

body link_from example
{
when_linking_children => "if_no_such_file";
}

```

Notes:

The options refer to what happens if the directory exists already and is already partially populated with files. If the directory being copied from contains a file with the same name as that of a link to be created, we must decide whether to override the existing destination object with a link or simply omit the automatic linkage for files that already

exist. The latter case can be used to make a copy of one directory with certain fields overridden.

7.2.4 perms (compound body)

Type: (ext body)

'mode' **Type:** string
Allowed input range: [0-7augorwxst,+~]+
Synopsis: File permissions (like posix chmod)
Example:

```
body perms example
{
mode => "a+rx,o+w";
}
```

Notes:

The mode string may be symbolic or numerical, like `chmod`.

'owners' **Type:** slist
Allowed input range: [a-zA-Z0-9_\$.~]+
Synopsis: List of acceptable owners or user ids, first is change target
Example:

```
body perms example
{
owners => { "mark", "wwwrun", "jeang" };
}
```

Notes:

The first user is the reference value that cfengine will set the file to if none of the list items matches the true state of the file.

'groups' **Type:** slist
Allowed input range: [a-zA-Z0-9_\$.~]+
Synopsis: List of acceptable groups of group ids, first is change target

Example:

```
body perms example
{
groups => { "users", "administrators" };
}
```

Notes:

The first named group in the list is the default that will be configured if the file does not match an element of the list.

'rxdirs' **Type:** (menu option)

Allowed input range:

```
    true
    false
    yes
    no
    on
    off
```

Synopsis: true/false add execute flag for directories if read flag is set

Example:

```
body perms rxdirs
{
rxdirs => "false";
}
```

Notes:

Default behaviour is to set the 'x' flag on directories automatically if the 'r' flag is specified when specifying multiple files in a single promise.

'bsdflags' **Type:** (option list)

Allowed input range:

```
    arch
    archived
    dump
```

```

opaque
sappnd
sappend
schg
schange
simmutable
sunlnk
sunlink
uappnd
uappend
uchg
uchange
uimmutable
uunlnk
unlink

```

Synopsis: List of menu options for bsd file system flags to set

Example:

body perms example

```

{
#..
bsdflags => { "uappnd", "uchg", "uunlnk", "nodump",
              "opaque", "sappnd", "schg", "sunlnk" };
}

```

Notes:

The free BSD Unices and MacOSX have additional filesystem flags which can be set. Refer to the BSD `chflags` documentation for this.

7.2.5 changes (compound body)

Type: (ext body)

'hash' **Type:** (menu option)

Allowed input range:

```

md5
sha1
best

```

Synopsis: Hash files for change detection

Example:

```
body changes example
{
hash => "md5";
}
```

Notes:

The **best** option cross correlates the best two available algorithms known in the OpenSSL library.

'report_changes'

Type: (menu option)

Allowed input range:

```
content
none
```

Synopsis: Specify criteria for change warnings

Example:

```
body changes example
{
report_changes => "content";
}
```

Notes:

Files can change in permissions and contents, i.e. external or internal attributes.

'update'

Type: (menu option)

Allowed input range:

```
true
false
yes
no
on
off
```

Synopsis: Update hash values immediately after change warning

Example:

```
body changes example
{
update => "true";
}
```

Notes:

File hashes should be updated as soon as a change is registered so that multiple warnings are not given about a single change.

7.2.6 delete (compound body)

Type: (ext body)

'dirlinks' **Type:** (menu option)

Allowed input range:

```
delete
tidy
keep
```

Synopsis: Menu option policy for dealing with symbolic links to directories during deletion

Example:

```
body delete example
{
dirlinks => "keep";
}
```

Notes:

Links to directories are normally removed just like any other link or file objects. By keeping directory links, you preserve the logical directory structure of the file system.

'rmdirs' **Type:** (menu option)

Allowed input range:

```

true
false
yes
no
on
off

```

Synopsis: true/false whether to delete empty directories during recursive deletion

Example:

```

body delete example
{
rmdir => "true";
}

```

Notes:

When deleting files recursively from a base directory, should we delete empty directories also, or keep the directory structure intact?

7.2.7 rename (compound body)

Type: (ext body)

'newname' **Type:** string

Allowed input range: (arbitrary string)

Synopsis: The desired name for the current file

Example:

```

body rename example(s)
{
newname => "${s}";
}

```

Notes:

'disable_suffix'

Type: string

Allowed input range: (arbitrary string)

Synopsis: The suffix to add to files when disabling (.cfdisabled)

Example:

```
body rename example
{
disable => "true";
disable_suffix => ".nuked";
}
```

Notes:

To make disabled files in a particular manner, use this string suffix. The default value is `‘.cf-disabled’`.

‘disable’ **Type:** (menu option)

Allowed input range:

```
true
false
yes
no
on
off
```

Synopsis: true/false automatically rename and remove permissions

Example:

```
body rename example
{
disable => "true";
disable_suffix => ".nuked";
}
```

Notes:

Disabling a file means making it impotent in the context in which it has an effect. For executables this means preventing execution, for an information file it means making the file unreadable.

‘rotate’ **Type:** int

Allowed input range: 0,99

Synopsis: Maximum number of file rotations to keep

Example:

```
body rename example
{
rotate => "4";
}
```

Notes:

Used for log rotation.

`'disable_mode'`

Type: string

Allowed input range: [0-7augorwxst,+~]+

Synopsis: The permissions to set when a file is disabled

Example:

```
body rename example
{
disable_mode => "0600";
}
```

Notes:

To disable an executable it is not enough to rename it, you should also remove the executable flag.

7.2.8 repository

Type: string

Allowed input range: [/\].*

Synopsis: Name of a repository for versioning

Example:

files:

```
"/path/file"
```



```
copy_from => source,
repository => "/var/cfengine/repository";
```

Notes:

A local repository for this object, overrides the default.

7.2.9 edit_line

Type: (ext bundle) (Separate Bundle)

7.2.10 edit_xml

Type: (ext bundle) (Separate Bundle)

7.2.11 edit_defaults (compound body)

Type: (ext body)

'edit_backup'

Type: (menu option)

Allowed input range:

```
true
false
timestamp
rotate
```

Synopsis: Menu option for backup policy on edit changes

Example:

```
body edit_defaults example
{
edit_backup => "timestamp";
}
```

Notes:

'max_file_size'

Type: int

Allowed input range: 0,99999999999

Synopsis: Do not edit files bigger than this number of bytes

Example:

```
body edit_defaults example
{
max_file_size => "50K";
}
```

Notes:

A local, per-file sanity check to make sure the file editing is sensible.

`'empty_file_before_editing'`

Type: (menu option)

Allowed input range:

```

true
false
yes
no
on
off
```

Synopsis: Baseline memory model of file to zero/empty before commencing promised edits

Example:

```
body edit_defaults example
{
empty_file_before_editing => "true";
}
```

Notes:

Emptying a file before reconstructing its contents according to a fixed recipe allows an ordered procedure to be convergent.

7.2.12 depth_search (compound body)

Type: (ext body)

`'include_dirs'`

Type: slist

Allowed input range: .*

Synopsis: List of regexes of directory names to include in depth search

Example:

```
body depth_search example
{
include_dirs => { "subdir1", "subdir2", "pattern.*" };
}
```

Notes:

This is the complement of `exclude_dirs`.

`'exclude_dirs'`

Type: slist

Allowed input range: .*

Synopsis: List of regexes of directory names NOT to include in depth search

Example:

```
body depth_search
{
# no dot directories
exclude_dirs => { "\..*" };
}
```

Notes:

Directory names are treated specially when searching recursively through a file system.

`'include_basedir'`

Type: (menu option)

Allowed input range:

```
true
false
yes
no
on
off
```

Synopsis: true/false include the start/root dir of the search results

Example:

```
body depth_search example
{
include_basedir => "true";
}
```

Notes:

When checking files recursively (with `depth_search`) the promiser is a directory. This parameter determines whether that initial directory should be considered part of the promise or simply a boundary which marks the edge of the search. If true, the promiser directory will also promise the same attributes as the files inside it.

‘depth’

Type: int**Allowed input range:** 0,99999999999**Synopsis:** Maximum depth level for search**Example:**

```
body depth_search example
{
depth => "inf";
}
```

Notes:

This was previous called ‘recurse’ in earlier versions of cfengine. Note that the value ‘inf’ may be used for an unlimited value.

‘xdev’

Type: (menu option)**Allowed input range:**

```
true
false
yes
no
on
off
```

Synopsis: true/false exclude directories that are on different devices

Example:

```
body depth_search example
{
xdev => "true";
}
```

Notes:

'traverse_links'

Type: (menu option)

Allowed input range:

```
    true
    false
    yes
    no
    on
    off
```

Synopsis: true/false traverse symbolic links to directories (false)

Example:

```
body depth_search example
{
traverse_links => "true";
}
```

Notes:

If this is true, `cf-agent` will treat symbolic links to directories as if they were directories. Normally this is considered a potentially dangerous assumption and links are not traversed.

'rmdeadlinks'

Type: (menu option)

Allowed input range:

```
    true
    false
```

yes
no
on
off

Synopsis: true/false remove links that point to nowhere

Example:

```
body depth_search example
{
  rmdeadlinks => "true";
}
```

Notes:

If we find links that point to non-existence files, should we delete them or keep them?

7.2.13 touch

Type: (menu option)

Allowed input range:

true
false
yes
no
on
off

Synopsis: true/false whether to touch time stamps on file

Example:

files:

```
"/path/file"
```

```
touch => "true";
```

Notes:

7.2.14 create

Type: (menu option)

Allowed input range:

```
true
false
yes
no
on
off
```

Synopsis: true/false whether to create non-existing file

Example:

files:

```
"/path/plain_file"
    create => "true";

"/path/dir/."
    create => "true";
```

Notes:

Directories are created by using the './' to signify a directory type. Note that, if no permissions are specified, mode 600 is chosen for a file, and mode 755 is chosen for a directory. If you cannot accept these defaults, you *should* specify permissions.

7.2.15 move_obstructions

Type: (menu option)

Allowed input range:

```
true
false
yes
no
on
off
```

Synopsis: true/false whether to move obstructions to file-object creation

Example:

files:

```

"/tmp/testcopy"

  copy_from    => mycopy("/tmp/source"),
  move_obstructions => "true",
  depth_search => recurse("inf");

```

Notes:

If we have promised to make file 'X' a link, but it already exists as a file, or vice-versa, or if a file is blocking the creation of a directory etc, then normally cfengine will report an error. If this is set, existing objects will be moved aside to allow the system to heal without intervention. Files and directories are saved/renamed, but symbolic links are deleted.

7.2.16 transformer**Type:** string**Allowed input range:** [/\].***Synopsis:** Shell command (with full path) used to transform current file**Example:**

```

"/home/mark/tmp/testcopy"

  file_select => pdf_files,
  transformer => "/usr/bin/gzip $(this.promiser),
  depth_search => recurse("inf");

```

Notes:

A command to execute on finding a file.

7.2.17 pathtype**Type:** (menu option)**Allowed input range:**

```

literal
regex

```


Synopsis: Menu option for interpreting promiser file object

Example:

```
files:

  "/var././lib"

  pathtype => "regex", #default
  perms => system;

  "/var././lib"

  pathtype => "literal",
  perms => system;
```

Notes:

If the keyword `literal` is invoked, a path looking like a regular expression will be treated as a literal string. Thus in the example, one case implies an iteration over all files/directories matching the regular expression, while the other means a single literal object with a name composed of dots and stars.

7.2.18 `acl` (compound body)

Type: (ext body)

`'acl_method'`

Type: (menu option)

Allowed input range:

```
append
prepend
set
```

Synopsis: Editing method for access control list

Example:

Fill me in (/home/mark/LapTop/CfengineProjects/CompanyDocuments/trunk/Cfengine3Reference/bodypr

Notes:

Fill me in (/home/mark/LapTop/CfengineProjects/CompanyDocuments/trunk/Cfengine3Reference/bodypr

`'acl_type'` **Type:** (menu option)

Allowed input range:

```
solaris
linux
ntfs
afs
```

Synopsis: Access control list type for the affected file system

Example:

Fill me in (/home/mark/LapTop/CfengineProjects/CompanyDocuments/trunk/Cfengine3Reference/body

Notes:

Fill me in (/home/mark/LapTop/CfengineProjects/CompanyDocuments/trunk/Cfengine3Reference/body

'acl_entry'

Type: slist

Allowed input range: .*

Synopsis: Native settings for access control entry

Example:

Fill me in (/home/mark/LapTop/CfengineProjects/CompanyDocuments/trunk/Cfengine3Reference/body

Notes:

Fill me in (/home/mark/LapTop/CfengineProjects/CompanyDocuments/trunk/Cfengine3Reference/body

7.3 insert_lines promises in 'edit_line'

This promise is part of the line-editing model. It inserts lines into the file at a specified location. The location is determined by body-attributes. The promise object referred to can be a literal line of a file-reference from which to read lines.

```
insert_lines:
    "literal line or file reference"
    location => location_body,
    ...;
```

```

body common control

{
any::

    bundlesequence => {
        example
    };
}

#####

bundle agent example

{
files:

    "/var/spool/cron/crontabs/root"

    edit_line => addline;
}

#####
# For the library
#####

bundle edit_line addline

{
insert_lines:

    "0,5,10,15,20,25,30,35,40,45,50,55 * * * * /var/cfengine/bin/cf-execd -F";
}

```

By parameterizing the editing bundle, one can make generic and reusable editing bundles.

7.3.1 location (compound body)

Type: (ext body)

'select_line_matching'

Type: string

Allowed input range: .*

Synopsis: Regular expression for matching file line location

Example:

```
body location example
{
select_line_matching => "^Expression match.* whole line$";
}
```

Notes:

The '^\$' are not necessary, just remember that the expression must match a whole line, not a fragment within a line.

'before_after'

Type: (menu option)

Allowed input range:

```
before
after
```

Synopsis: Menu option, point cursor before of after matched line

Example:

```
body location append
{
#...
before_after => "before";
}
```

Notes:

Determines whether an edit will occur before or after the currently matched line.

'first_last'

Type: (menu option)

Allowed input range:

```
first
last
```

Synopsis: Menu option, choose first or last occurrence of match in file

Example:

```
body location example
{
first_last => "last";
}
```

Notes:

In multiple matches, decide whether the first or last occurrence of the matching pattern in the case affected by the change. In principle this could be generalized to more cases but this seems like a fragile quality to evaluate, and only these two cases are deemed of reproducible significance.

7.3.2 insert_type

Type: (menu option)

Allowed input range:

```
literal
string
file
```

Synopsis: Type of object the promiser string refers to (default literal)

Example:

```
body insert_lines example
{
insert_type => "file";
}
```

Notes:

The default is to treat the promiser as a literal string. This is used to tell cfengine that the string is non-literal and should be interpreted as a filename from which to import lines.

7.3.3 insert_select (compound body)

Type: (ext body)

'insert_if_startwith_from_list'

Type: slist

Allowed input range: .*

Synopsis: Insert line if it starts with a string in the list

Example:

```
body insert_select example
{
insert_if_startwith_from_list => { "find_me_1", "find_me_2" };
}
```

Notes:

The list contains literal strings to search for in an secondary file (not the main file being edited). If the string is found as the first characters (at the start) of a line in the file, that line from the secondary file will be inserted at the present location in the primary file.

`'insert_if_not_startwith_from_list'`

Type: slist

Allowed input range: .*

Synopsis: Insert line if it DOES NOT start with a string in the list

Example:

```
body insert_select example
{
insert_if_not_startwith_from_list => { "find_me_1", "find_me_2" };
}
```

Notes:

The complement of `insert_if_startwith_from_list`.

`'insert_if_match_from_list'`

Type: slist

Allowed input range: .*

Synopsis: Insert line if it fully matches a regex in the list

Example:

```
body insert_select example
```

```
{
insert_if_match_from_list => { ".*find_.*_1.*", ".*find_.*_2.*" };
}
```

Notes:

The list contains regular expressions to search for in an secondary file (not the main file being edited). If the regex matches a complete line of the file, that line from the secondary file will be inserted at the present location in the primary file.

```
'insert_if_not_match_from_list'
```

Type: slist

Allowed input range: .*Insert line if it DOES NOT fully match a regex in the list

Synopsis: (null)

Example:

```
body insert_select example
{
insert_if_not_match_from_list => { ".*find_.*_1.*", ".*find_.*_2.*" };
}
```

Notes:

The complement of `insert_if_match_from_list`.

```
'insert_if_contains_from_list'
```

Type: slist

Allowed input range: .*

Synopsis: Insert line if a regex in the list match a line fragment

Example:

```
body insert_select example
{
insert_if_contains_from_list => { "find_me_1", "find_me_2" };
}
```

Notes:

The list contains literal strings to search for in an secondary file (not the main file being edited). If the string is found in a line of the file, that line from the secondary file will be inserted at the present location in the primary file.

`'insert_if_not_contains_from_list'`

Type: slist

Allowed input range: .*

Synopsis: Insert line if a regex in the list DOES NOT match a line fragment

Example:

```
body insert_select example
{
insert_if_not_contains_from_list => { "find_me_1", "find_me_2" };
}
```

Notes:

The complement of `insert_if_contains_from_list`.

7.3.4 `expand_scalars`

Type: (menu option)

Allowed input range:

```
true
false
yes
no
on
off
```

Synopsis: Expand any unexpanded variables

Example:

```
body insert_lines example
{
insert_type    => "file";
expand_scalars => "true";
}
```

Notes:

A way of incorporating templates with variable expansion into file operations. Variables should be named and scoped appropriately for the bundle in which this promise is made.

In cfengine 2 `editfiles` this was called 'ExpandVariables'.

7.4 `field_edits` promises in 'edit_line'

Certain types of text file (e.g. the 'passwd' and 'group' files in Unix) are tabular in nature, with field separators (e.g. ':' or ','). This promise assumes a parameterizable model for editing the fields of such files, using a regular expression to separate major fields and a character to separate sub-fields. First you match the line with a regular expression, then a `field_edits` body describes the separators for fields and one level of sub-fields, along with policies for editing these fields, ordering the items within them etc.

```
field_edits:
    "regex matching line"
        edit_field => body;
```

bundle agent example

```
{
vars:

    "userset" slist => { "one-x", "two-x", "three-x" };

files:

    "/tmp/passwd"

        create      => "true",
        edit_line => SetUserParam("mark","6","/set/this/shell");

    "/tmp/group"

        create      => "true",
        edit_line => AppendUserParam("root","4","@(userset)");
}
```

```
#####
```

```
bundle edit_line SetUserParam(user,field,val)
```

```

{
field_edits:

    "$(user).*"

    # Set field of the file to parameter

    edit_field => col(":", "$(field)", "$(val)", "set");
}

#####

bundle edit_line AppendUserParam(user,field,allusers)
{
vars:

    "val" slist => { @(allusers) };

field_edits:

    "$(user).*"

    # Set field of the file to parameter

    edit_field => col(":", "$(field)", "$(val)", "alphanum");

}

#####
# Bodies
#####

body edit_field col(split,col,newval,method)

{
field_separator => "$(split)";
select_field    => "$(col)";
value_separator => ",";
field_value     => "$(newval)";
field_operation => "$(method)";
extend_fields   => "true";
}

```

Field editing allows us to edit tabular files in a unique way, adding and removing data from addressable fields. The 'passwd' and 'group' files are classic examples of tabular files, but there are many ways to use this feature, e.g. edit a string

```
VARIABLE="one two three"
```

View this line as a tabular line separated by '"' and with sub-separator given by the space.

7.4.1 edit_field (compound body)

Type: (ext body)

'field_separator'

Type: string

Allowed input range: .*

Synopsis: The regular expression used to separate fields in a line

Example:

```
body edit_field example
{
field_separator =>> ":";
}
```

Notes:

Most tabular files are separated by simple characters, but by allowing a general regular expression one can make creative use of this model to edit all kinds of line-based text files.

'select_field'

Type: int

Allowed input range: 0,9999999999

Synopsis: Integer index of the field required 1..n

Example:

```
body field_edits example
{
select_field => "5";
}
```

Notes:

Numbering starts from 1 not from 0.

`'value_separator'`

Type: string

Allowed input range: `^.$`

Synopsis: Character separator for subfields inside the selected field

Example:

```
body field_edit example
{
value_separator => ",";
}
```

Notes:

For example, elements in the group file are separated by ':', but the lists of users in these fields are separated by ','.

`'field_value'`

Type: string

Allowed input range: `.*`

Synopsis: Set field value to a fixed value

Example:

```
body edit_field example(s)
{
field_value => "$(s)";
}
```

Notes:

Set a field to a constant value, e.g. reset the value to a constant default, empty the field, or set it fixed list.

`'field_operation'`

Type: (menu option)

Allowed input range:

prepend

```

append
alphanum
delete
set

```

Synopsis: Menu option policy for editing subfields

Example:

```

body edit_field example
{
field_operation => "append";
}

```

Notes:

The method by which to edit a field in multi-field/column editing of tabular files.

'extend_fields'

Type: (menu option)

Allowed input range:

```

true
false
yes
no
on
off

```

Synopsis: true/false add new fields at end of line if necessary to complete edit

Example:

```

body edit_field example
{
extend_fields => "true";
}

```

Notes:

Blank fields in a tabular file can be eliminated or kept depending in this setting. If in doubt, set this to true.

`'allow_blank_fields'`

Type: (menu option)

Allowed input range:

```

true
false
yes
no
on
off

```

Synopsis: true/false allow blank fields in a line (do not purge)

Example:

```

body edit_field example
{
# ...
allow_blank_fields => "true";
}

```

Notes:

When editing a file using the field or column model, blank fields, especially at the start and end are generally discarded. If this is set to true, cfengine will retain the blank fields and print the appropriate number of field separators.

7.5 `replace_patterns` promises in `'edit_line'`

This promise refers to arbitrary text patterns in a file. The pattern is expressed as a regular expression and must be compatible with the default model for regular expressions on your system. The default model is PCRE (Perl Compatible Regular Expressions) if available.

```

replace_patterns:
  "search pattern"
  replace_with => replace_body,
  ...;

```

Fill me in (/home/mark/LapTop/Cfengine3/trunk/docs/promise_replace_patterns_notes.texinfo)■

7.5.1 `replace_with` (compound body)

Type: (ext body)

`'replace_value'`

Type: string

Allowed input range: .*

Synopsis: Value used to replace regular expression matches in search

Example:

```
body replace_with example(s)
{
  replace_value => "$(s)";
}
```

Notes:

`'occurrences'`

Type: (menu option)

Allowed input range:

```
all
first
```

Synopsis: Menu option to replace all occurrences or just first

Example:

```
body replace_with example
{
  occurrences => "first";
}
```

Notes:

A policy for string replacement.

7.6 delete_lines promises in 'edit_line'

This promise assures that certain lines matching regular expression patterns exactly will not be present in a text file. If the lines are found, the default promise is to remove them.

```
bundle edit_line example
{
  delete_lines:

    "olduser.*";

}
```

7.6.1 not_matching

Type: (menu option)

Allowed input range:

```
    true
    false
    yes
    no
    on
    off
```

Synopsis: true/false negate match criterion

Example:

```
delete_lines:

  # edit /etc/passwd

  "mark.*|root.*" not_matching => "true";
```

Notes:

The negation of an expression (for convenience).

7.6.2 delete_select (compound body)**Type:** (ext body)`'delete_if_startwith_from_list'`**Type:** slist**Allowed input range:** .***Synopsis:** Delete line if it starts with a string in the list**Example:**

```
body delete_select example(s)
{
delete_if_startwith_from_list => { @(s) };
}
```

Notes:

Delete lines from a file if they begin with the sub-strings listed.

`'delete_if_not_startwith_from_list'`**Type:** slist**Allowed input range:** .***Synopsis:** Delete line if it DOES NOT start with a string in the list**Example:**

```
body delete_select example(s)
{
delete_if_not_startwith_from_list => { @(s) };
}
```

Notes:

Delete lines from a file unless they start with the sub-strings in the list given.

`'delete_if_match_from_list'`**Type:** slist**Allowed input range:** .***Synopsis:** Delete line if it fully matches a regex in the list**Example:**

```
body delete_select example(s)
{
delete_if_match_from_list => { @(s) };
}
```

Notes:

Delete lines from a file if they completely match the regular expressions listed.

`'delete_if_not_match_from_list'`

Type: slist

Allowed input range: .*

Synopsis: Delete line if it DOES NOT fully match a regex in the list

Example:

```
body delete_select example(s)
{
delete_if_not_match_from_list => { @(s) };
}
```

Notes:

Delete lines from a file unless they fully match regular expressions in the list.

`'delete_if_contains_from_list'`

Type: slist

Allowed input range: .*

Synopsis: Delete line if a regex in the list match a line fragment

Example:

```
body delete_select example(s)
{
delete_if_contains_from_list => { @(s) };
}
```

Notes:

Delete lines from a file if they contain the sub-strings listed.

`'delete_if_not_contains_from_list'`

Type: slist

Allowed input range: .*

Synopsis: Delete line if a regex in the list DOES NOT match a line fragment

Example:

```
body delete_select discard(s)
{
delete_if_not_contains_from_list => { "substring1", "substring2" };
}
```

Notes:

Delete lines from the file which do not contain the sub-strings listed.

```
bundle agent main(parameter)
{
vars:

"sys_files"    slist        => {
                        "/etc/passwd",
                        "/etc/services"
                    };

files:

"${sys_files}" perms      => p("root", "0644"),
changes        => trip_wire;

"/etc/shadow"  perms      => p("root", "0600"),
changes        => trip_wire;

"/usr"         changes    => trip_wire,
depth_search  => recurse("inf");

"/tmp"         delete     => tidy,
file_select   => days("2"),
depth_search  => recurse("inf");
}
```

Agent bundles contain user-defined promises for `cf-agent`. The types of promises and their corresponding bodies are detailed below.

7.7 interfaces promises in ‘agent’

Interfaces promises describe the configurable aspects relating to network interfaces. Most workstations and servers have only a single network interface, but routers and multi-homed hosts often have multiple interfaces. Interface promises include attributes such as the IP address identity, assumed netmask and routing policy in the case of multi-homed hosts. For virtual machines and hosts, the list of interfaces can be quite large.

```
interfaces:
  "interface name"
    tcp_ip => tcp_ip_body,
    ...;
```

Fill me in (/home/mark/LapTop/Cfengine3/trunk/docs/promise_interfaces_example.texinfo)

Fill me in (/home/mark/LapTop/Cfengine3/trunk/docs/promise_interfaces_notes.texinfo)

7.7.1 tcp_ip (compound body)

Type: (ext body)

‘ipv4_address’

Type: string

Allowed input range: [0-9.]+/[0-4]+

Synopsis: IPv4 address for the interface

Example:

```
body tcp_ip example
{
  ipv4_address => "123.456.789.001";
}
```

Notes:

The address will be checked and if necessary set. Today few hosts will be managed in this way: address management will be handled by other services like DHCP.

`'ipv4_netmask'`

Type: string

Allowed input range: [0-9.]+/[0-4]+

Synopsis: Netmask for the interface

Example:

```
body tcp_ip example
{
  ipv4_netmask => "255.255.254.0";
}
```

Notes:

In many cases the CIDR form of address will show the netmask as '/23', but this offers and 'old style' alternative.

`'ipv6_address'`

Type: string

Allowed input range: [0-9a-fA-F:]+/[0-9]+

Synopsis: IPv6 address for the interface

Example:

```
body tcp_ip example
{
  ipv6_address => "2001:700:700:3:20f:1fff:fe92:2cd3/64";
}
```

Notes:

```

bundle agent main(parameter)
{
vars:

  "sys_files"    slist      => {
                        "/etc/passwd",
                        "/etc/services"
                    };

files:

  "${sys_files}" perms      => p("root","0644"),
                 changes    => trip_wire;

  "/etc/shadow"  perms      => p("root","0600"),
                 changes    => trip_wire;

  "/usr"         changes    => trip_wire,
                 depth_search => recurse("inf");

  "/tmp"         delete     => tidy,
                 file_select => days("2"),
                 depth_search => recurse("inf");
}

```

Agent bundles contain user-defined promises for `cf-agent`. The types of promises and their corresponding bodies are detailed below.

7.8 methods promises in ‘agent’

Methods are compound promises that refer to whole bundles of promises. Methods may be parameterized. Methods promises are written in a form that is ready for future development. The promiser object is an abstract identifier that refers to a collection (or pattern) of lower level objects that are affected by the promise-bundle. Since the use of these identifiers is for the future, you can simply use any string here for the time being.

```

methods:

  "any"

    usebundle => method_id("parameter",...);

```

Methods are useful for encapsulating repeatedly used configuration issues and iterating over parameters.

In cfengine 2 methods referred to separate sub-programs executed as separate processes. Methods are now implemented as bundles that are run inline.

```

bundle agent example
{
vars:

  "userlist" slist => { "mark", "jeang", "jonhenrik", "thomas", "eben" };

methods:

  "any" usebundle => subtest("${userlist}");
}

#####

bundle agent subtest(user)

{
commands:

  "/bin/echo Fix ${user}";

reports:

  linux::

    "Finished doing stuff for ${user}";
}

```

Methods offer powerful ways to encapsulate multiple issues pertaining to a set of parameters.

7.8.1 usebundle

Type: (ext bundle) (Separate Bundle)

```

bundle agent main(parameter)
{
vars:

  "sys_files"    slist      => {
                        "/etc/passwd",
                        "/etc/services"
                    };

files:

  "${sys_files}" perms      => p("root","0644"),
                  changes   => trip_wire;

  "/etc/shadow"  perms      => p("root","0600"),
                  changes   => trip_wire;

  "/usr"         changes    => trip_wire,
                  depth_search => recurse("inf");

  "/tmp"         delete     => tidy,
                  file_select => days("2"),
                  depth_search => recurse("inf");
}

```

Agent bundles contain user-defined promises for `cf-agent`. The types of promises and their corresponding bodies are detailed below.

7.9 packages promises in ‘agent’

Package promises refer to native packages provided by operating system releases. Package management is rapidly developing area and this feature is currently undergoing further research.

```

packages:

  ?

```

Users may continue to use the package management features in cfengine 2 until the final design for package promises is finalized.

Fill me in (/home/mark/LapTop/Cfengine3/trunk/docs/promise_packages_example.texinfo)

Fill me in (/home/mark/LapTop/Cfengine3/trunk/docs/promise_packages_notes.texinfo)

7.9.1 install (compound body)

Type: (ext body)

```

bundle agent main(parameter)
{
vars:

  "sys_files"    slist      => {
                        "/etc/passwd",
                        "/etc/services"
                      };

files:

  "${sys_files}" perms      => p("root","0644"),
                  changes   => trip_wire;

  "/etc/shadow"  perms      => p("root","0600"),
                  changes   => trip_wire;

  "/usr"         changes    => trip_wire,
                  depth_search => recurse("inf");

  "/tmp"         delete     => tidy,
                  file_select => days("2"),
                  depth_search => recurse("inf");
}

```

Agent bundles contain user-defined promises for `cf-agent`. The types of promises and their corresponding bodies are detailed below.

7.10 processes promises in ‘agent’

Process promises refer to items in the system process table. Note that this is not the same as commands (which are instructions). A process is a command in some state of execution (with a Process Control Block). Promiser objects here are patterns that match line fragments in the system process table.

```

processes:

  "regex contained in process line"

    process_select => process_filter_body,
    restart_class => "activation class for process",
    ..;

```

In cfengine 2 there was a restart clause for directly executing a command to restart a process. In cfengine 3 there is instead a class to activate. You must then describe a `command` in that class to restart the process.

```

commands:

```

```

  restart_me::

    "/path/executable" ... ;

```

This rationalizes complex restart-commands and avoids unnecessary overlap between `processes` and `commands`.

The `process_stop` is also arguably a command, but it should be an ephemeral command that does not lead to a persistent process. It is intended only for commands of the form ‘`/etc/inetd service stop`’, not for processes that persist. Processes are restarted at the end of a bundle’s execution, but stop commands are executed immediately.

Take care to note that process table formats differ between operating systems, and the use of simple patterns such as program-names is recommended. For more sophisticated matches, users should use the `process_select` feature.

Note: `process_select` was previously called `process filters` in cfengine 2 and earlier.

```

bundle agent example
{
processes:

  ".*"

    process_count    => anyprocs,
    process_select   => proc_finder;

reports:

  any_procs::

    "Found processes out of range";
}

```

```
#####

body process_select proc_finder

{
stime_range => irange(ago(0,0,0,5,30,0),ago(0,0,0,0,20,0));
process_result => "stime";
}

#####

body process_count anyprocs

{
match_range => "0,0";
out_of_range_define => { "any_procs" };
}

```

In cfengine 2, one has two separate actions:

```
processes
shellcommands
```

In cfengine 3 we have

```
processes
commands
```

Cfengine 2 got this ontology about right intuitively, but not quite. It allowed a `restart` command to appear in a process promise, which is really a command execution. This has been changed in cfengine 3 so that there is a cleaner separation. Let's see why.

Executions are about jobs, services, scripts etc. They are properties of an executable file. The referring 'promiser' is a file object. On the other hand a process is a property of a "process identifier" which is a kernel instantiation, a quite different object altogether. So it makes sense to say that

- A "PID" (which is not an executable) promises to be reminded of a signal, e.g.

```
kill signal pid
```
- An "command" promises to start or stop itself with a parameterized specification.

```
exec command argument1 argument2 ...
```

Neither the file nor the pid necessarily promise to respond to these activations, but they are nonetheless physically meaningful phenomena or attributes associated with these objects.

- Executable files do not listen for signals as they have no active state.
- PIDs do not run themselves or stop themselves with new arguments, but they can use signals as they are running.

Executions lead to processes for the duration of their lifetime, so these two issues are related, although the promises themselves are not.

Services versus processes:

A service is an abstraction that requires processes to run and files to be configured. It makes a lot of sense to wrap services in modular bundles. Starting and stopping a service can be handled in at least two ways. Take the web service as an example.

We can start the service by promising an execution of a daemon (e.g. `httpd`). Normally this execution does not terminate without intervention. We can terminate it in one of two ways:

- Using a process signal, by promising a signal to processes matching a certain pid search
- Using an execution of a termination command, e.g. `'/etc/init.d/apache stop'`.

The first case makes sense if we need to qualify the termination by searching for the processes. The processes section of a cfengine 3 policy includes a control promise to search for matching processes. If matches are found, signals can be sent to precisely each specific process.

Classes can also be defined, in principle triggering an execution of the stop script, but then the class refers only to the presence of matching pids, not to the individual pids concerned. So it becomes the responsibility of the execution to locate and interact with the pids necessary.

Want it running?:

How do we say simply that we want a service running? In the agent control promises, we could check each service individually.

```
bundlesequence => { Update, Service("apache"), Service("nfsd") };
```

or

```
bundlesequence => { Update, @(globals.all_services) };
```

The bundle for this can look like this:

```
bundle agent Service("${service}")

{
processes:

    "${service}"

        process_count => up("${service}");

commands:

    "${daemons[${service}]}"

        ifvarclass => "${service}_up",
        args       => "${args[${service}]}";

}
```

An alternative would be self-contained:

```
bundle agent Service

{
vars:
```

```

"service" slist => { "apache", "nfsd", "bind" };

processes:

"$(service)"

    process_count => up("$(service)");

commands:

"$daemons[$(service)]"

    ifvarclass => "$(service)_up",
    args       => "$args[$(service)]";

}

#####
# Parameterized body
#####

body process_count("$(s)")

{
match_range => "[0,10]";
out_of_range_define => "$(s)_up";
}

```

Is this a step backwards? The cfengine 3 approach might seem like a step backwards from the simple cfengine 2 statement:

```

processes:

"httpd" restart "/etc/init.d/apache restart"

```

However, it allows several improvements.

You can do other things in between stopping and starting the service, like file editing, or security sweeps. You can use templates to simplify the syntax in bulk for several process checks or restarts.

```

processes:

"$(service.list)"

```

If you don't want any delay in stopping and starting the service, then place these promises in a private bundle with nothing in between them.

7.10.1 signals

Type: (option list)

Allowed input range:

```

    hup
    int
    trap
    kill
    pipe
    cont
    abrt
    stop
    quit
    term
    child
    usr1
    usr2
    bus
    segv

```

Synopsis: A list of menu options representing signals to be sent to a process

Example:

processes:

```

cfserverd_out_of_control::

    "cfserverd"

        signals      => { "stop" , "term" },
        restart_class => "start_cfserverd";

any::

    "snmpd"

        signals      => { "term" , "kill" };

```

Notes:

Signals are presented as an ordered list until the first one succeeds.

7.10.2 process_stop

Type: string

Allowed input range: [/\].*

Synopsis: A command used to stop a running process

Example:

processes:

```
"snmpd"

    process_stop => "/etc/init.d/snmp stop";
```

Notes:

As an alternative to sending a termination or kill signal to a process, one may call a ‘stop script’ to perform a graceful shutdown.

7.10.3 process_count (compound body)**Type:** (ext body)

‘match_range’

Type: irange [int,int]**Allowed input range:** 0,9999999999**Synopsis:** Integer range for acceptable number of matches for this process**Example:**

```
body process_count example
{
  match_range => irange("10","50");
}
```

Notes:

This is a numerical range for the number of occurrences of the process in the process table. As long as it falls within the specified limits, the promise is considered kept.

‘in_range_define’

Type: slist**Allowed input range:** (arbitrary string)**Synopsis:** List of classes to define if the matches are in range**Example:**

```
body process_count example
{
in_range_define => { "class1", "class2" };
}
```

Notes:

Classes are defined if the processes that are found in the process table satisfy the promised process count, i.e. if the promise about the number of processes matching the other criteria is kept.

`'out_of_range_define'`

Type: slist

Allowed input range: (arbitrary string)

Synopsis: List of classes to define if the matches are out of range

Example:

```
body process_count example(s)
{
out_of_range_define => { "process_anomaly", "anomaly_$(s)"};
}
```

Notes:

Classes to activate remedial promises conditional on this promise failure to be kept.

7.10.4 process_select (compound body)

Type: (ext body)

`'process_owner'`

Type: slist

Allowed input range: (arbitrary string)

Synopsis: List of regexes matching the user of a process

Example:

```
body process_select example
{
process_owner => { "wwwrun", "nobody" };
}
```



```
}

```

Notes:

Regular expression should match a legal user name on the system.

'pid'

Type: irange [int,int]

Allowed input range: 0,9999999999

Synopsis: Range of integers matching the process id of a process

Example:

```
body process_select example
{
  pid => irange("1","10");
  process_result => "pid";
}
```

Notes:

'ppid'

Type: irange [int,int]

Allowed input range: 0,9999999999

Synopsis: Range of integers matching the parent process id of a process

Example:

```
body process_select example
{
  ppid => irange("407","511");
}
```

Notes:

'pgid'

Type: irange [int,int]

Allowed input range: 0,9999999999

Synopsis: Range of integers matching the parent group id of a process

Example:

```
body process_select example
{
  pgid => irange("1","10");
  process_result => "pgid";
}
```

Notes:

'rsize' **Type:** irange [int,int]
Allowed input range: 0,9999999999
Synopsis: Range of integers matching the resident memory size of a process
Example:

```
body process_select
{
  rsize => irange("4000","8000");
}
```

Notes:

'vsize' **Type:** irange [int,int]
Allowed input range: 0,9999999999
Synopsis: Range of integers matching the virtual memory size of a process
Example:

```
body process_select example
{
  vsize => irange("4000","9000");
}
```

Notes:

'status' **Type:** string
Allowed input range: (arbitrary string)

Synopsis: Regular expression matching the status field of a process

Example:

```
body process_select example
{
status => "Z";
}
```

Notes:

For instance, characters in the set 'NRS<s1+..''

'ttime_range'

Type: irange [int,int]

Allowed input range: 0,4026531839

Synopsis: Range of integers matching the total elapsed time of a process

Example:

```
body process_select example
{
ttime_range => irange(0,accumulated(0,1,0,0,0,0));
}
```

Notes:

This is total accumulated time for a process.

'stime_range'

Type: irange [int,int]

Allowed input range: 0,4026531839

Synopsis: Range of integers matching the start time of a process

Example:

```
body process_select example
{
stime_range => irange(ago(0,0,0,1,0,0),now);
}
```

Notes:

The calculation of time from process table entries is sensitive to Daylight Savings Time (Summer/Winter Time) so calculations could be a hour off. This is for now a bug to be fixed.

'command'

Type: string**Allowed input range:** (arbitrary string)**Synopsis:** Regular expression matching the command/cmd field of a process**Example:**

```
body select_process example
{
command => "cf-.*";

process_result => "command";
}
```

Notes:

This expression should match the entire **COMMAND** field of the process table (not just a fragment). This field is usually the last field on the line and thus starts with the first non-space character and ends with the end of line.

'tty'

Type: string**Allowed input range:** (arbitrary string)**Synopsis:** Regular expression matching the tty field of a process**Example:**

```
body process_select example
{
tty => "pts/[0-9]+";
}
```

Notes:

`'priority'` **Type:** `irange [int,int]`
Allowed input range: `-20,+20`
Synopsis: Range of integers matching the priority field (PRI/NI) of a process
Example:

```
body process_select example
{
priority => irange("-5","0");
}
```

Notes:

`'threads'` **Type:** `irange [int,int]`
Allowed input range: `0,9999999999`
Synopsis: Range of integers matching the threads (NLWP) field of a process
Example:

```
body process_select example
{
threads => irange(1,5);
}
```

`@verbatim`

`@b{Notes}:@*`
`@*`

`@item @code{process_result}`
`@b{Type}: string`

`@noindent @b{Allowed input range}: @code{[(process_owner|pid|ppid|pgid|rsize|vsize|statu`

`@noindent @b{Synopsis}: Boolean class expression returning the logical combination of cla`

`@b{Example}:@*`

```

@*

@verbatim

body process_select proc_finder(p)

{
process_owner => { "avahi", "bin" };
command       => "$(p)";
pid           => "100,199";
vsize        => "0,1000";
process_result => "command.(process_owner|rsize)";
}

```

Notes:

A logical combination of the process selection classifiers. The syntax is the same as that for class expressions. There should be no spaces in the expressions.

7.10.5 restart_class

Type: string

Allowed input range: [a-zA-Z0-9_\$.]+

Synopsis: A class to be set if the process is not running, so that a command: rule can be referred to restart the process

Example:

```

processes:

    "cfservd"

    restart_class => "start_cfserv";

commands:

    start_cfserv::

        "/usr/local/sbin/cfservd";

```

Notes:

This is a signal to restart a process that should be running, if it is not running. Processes are signalled first and then restarted later, at the end of bundle execution, after all possible corrective actions have been made that could influence their execution.

```

bundle agent main(parameter)
{
vars:
  "sys_files"  slist      => {
                        "/etc/passwd",
                        "/etc/services"
                      };
files:
  "${sys_files}" perms    => p("root", "0644"),
                  changes => trip_wire;

  "/etc/shadow" perms    => p("root", "0600"),
                  changes => trip_wire;

  "/usr"         changes => trip_wire,
                  depth_search => recurse("inf");

  "/tmp"         delete   => tidy,
                  file_select => days("2"),
                  depth_search => recurse("inf");
}

```

Agent bundles contain user-defined promises for `cf-agent`. The types of promises and their corresponding bodies are detailed below.

7.11 storage promises in ‘agent’

Storage promises refer to disks and filesystem properties.

```

storage:
  "/disk volume or mountpoint"
    volume => volume_body,
    ...;

```

In cfengine 2, storage promises were divided into `disks` or `required`, and `misc_mounts` types. The old mount-models for binary and home servers has been deprecated and removed from cfengine 3. Users who use these models can reconstruct them from the low-level tools.

```
bundle agent storage

{
storage:

    "/usr" volume => mycheck("10%");
    "/mnt" mount  => nfs("nfsserv.example.org", "/home");

}

#####

body volume mycheck(free) # reusable template

{
check_foreign => "false";
freespace     => "$free";
sensible_size => "10000";
sensible_count => "2";
}

body mount nfs(server,source)

{
mount_type => "nfs";
mount_source => "$(source)";
mount_server => "$(server)";
editfstab => "true";
}
```

7.11.1 mount (compound body)

Type: (ext body)

`'mount_type'`

Type: (menu option)

Allowed input range:

```
nfs
nfs2
```



```
nfs3
nfs4
```

Synopsis: Protocol type of remote file system

Example:

```
body mount example
{
mount_type => "nfs3";
}
```

Notes:

This field is mainly for future extensions.

'mount_source'

Type: string

Allowed input range: [/\].*

Synopsis: Path of remote file system to mount

Example:

```
body mount example
{
mount_source "/location/disk/directory";
}
```

Notes:

This is the location on the remote device, server, SAN etc.

'mount_server'

Type: string

Allowed input range: (arbitrary string)

Synopsis: Hostname or IP or remote file system server

Example:

```
body mount example
{
```

```
mount_server => "nfs_host.example.org";
}
```

Notes:

Hostname or IP address, this could be on a SAN.

'mount_options'

Type: slist

Allowed input range: List of option strings to add to the file system table ("fstab")

Synopsis: (null)

Example:

```
body mount example
{
mount_options => { "rw", "acls" };
}
```

Notes:

This list is concatenated in a form appropriate for the filesystem. The options must be legal options for the system mount commands.

'edit_fstab'

Type: (menu option)

Allowed input range:

```

true
false
yes
no
on
off
```

Synopsis: true/false add or remove entries to the file system table ("fstab")

Example:

```
body mount example
{
edit_fstab => "true";
```

```
}

```

Notes:

The default behaviour is to not place edits in the file system table.

`'unmount'` **Type:** (menu option)

Allowed input range:

```

true
false
yes
no
on
off

```

Synopsis: true/false unmount a previously mounted filesystem

Example:

```

body mount example
{
unmount => "true";
}

```

Notes:

7.11.2 volume (compound body)

Type: (ext body)

`'check_foreign'`

Type: (menu option)

Allowed input range:

```

true
false
yes
no
on
off

```

Synopsis: true/false verify storage that is mounted from a foreign system on this host

Example:

```
body volume example
{
#..
check_foreign => "false";
}
```

Notes:

Cfengine will not normally perform sanity checks on filesystems which are not local to the host. If `true` it will ignore a partition's network location and ask the current host to verify storage located physically on other systems.

'freespace'

Type: string

Allowed input range: [0-9]+[mb%]

Synopsis: Absolute or percentage minimum disk space that should be available before warning

Example:

```
body volume example1
{
freespace => "10%";
}
```

```
body volume example2
{
freespace => "50M";
}
```

Notes:

The amount of freespace that is promised on a storage device. Once this promise is found not to be kept, warnings are generated.

'sensible_size'

Type: int

Allowed input range: 0,999999999999

Synopsis: Minimum size in bytes that should be used on a sensible-looking storage device

Example:

```
body volume example
{
sensible_size => "20K";
}
```

Notes:

```
body volume control
{
sensible_size => "20K";
}
```

'sensible_count'

Type: int

Allowed input range: 0,9999999999

Synopsis: Minimum number of files that should be defined on a sensible-looking storage device

Example:

```
body volume example
{
sensible_count => "20";
}
```

Notes:

Files must be readable by the agent, i.e. it is assumed that the agent has privileges on volumes being checked.

'scan_arrivals'

Type: (menu option)

Allowed input range:

```
true
false
yes
no
on
off
```

Synopsis: true/false generate pseudo-periodic disk change arrival distribution

Example:

Fill me in (/home/mark/LapTop/Cfengine3/trunk/docs/bodypart_scan_arrivals_example.texinfo)

```
body volume example
{
scan_arrivals => "true";
}
```

Notes:

This operation should not be left 'on' for more than a single run (maximum once per week). It causes cfengine to perform an extensive disk scan noting the schedule of changes between files. This can be used for a number of analyses including optimum backup schedule computation.

8 Bundles of server

```
bundle server access_rules()
{
access:

  "/home/mark/PrivateFiles"

  admit  => { ".*\.example\.org" };

  "/home/mark/.cfagent/bin/cf-agent"

  admit  => { ".*\.example\.org" };

roles:

  ".*"  authorize => { "mark" };
}
```

Bundles in the server describe access promises on specific file and class objects supplied by the server to clients.

8.1 access promises in ‘server’

Access promises are conditional promises made by the server about file objects. The promise has two consequences. For file copy requests, the file becomes transferrable to the remote client according to the conditions specified in the server promise (i.e. if the connection encryption requirements are met, and if the client has been granted appropriate privileges with `maproot` (like its NFS counterpart) to be able to see file objects not owned by the server process owner).

The promise has two mutually exclusive attributes ‘`admit`’ and ‘`deny`’. Use of ‘`admit`’ is preferred as mistakes and omissions can easily be made when excluding from a group.

When access is granted to a directory, the promise is automatically given about all of its contents and sub-directories. The access promise allows overlapping promises to be made, and these are kept in a first-come-first-served fashion. Thus file objects (promisers) should be listed in order of most-specific file first. In this way, specific promises will override less specific ones.

```

access:

  "/path/file_object"

  admit  => { "hostname", "ipv4_address", "ipv6_address" };

```

8.1.1 Access Example

```

#####
# Server config
#####

body server control

{
allowconnects      => { "127.0.0.1" , "::1" };
allowallconnects  => { "127.0.0.1" , "::1" };
trustkeysfrom     => { "127.0.0.1" , "::1" };
}

#####

bundle server access_rules()

{
access:

  "/home/mark/LapTop"

  admit  => { "127.0.0.1" };
}

```

Entries may be literal addresses of IPv4 or IPv6, or any name registered in the POSIX gethostbyname service.

8.1.2 admit

Type: slist

Allowed input range: (arbitrary string)

Synopsis: List of host names or IP addresses to grant access to file objects

Example:

```
access:
```

```
"/home/mark/LapTop"
```

```
admit => { "127.0.0.1", "192.168.0.1", "*.domain.tld" };
```

Notes:

Admit promises grant access to file objects on the server. Arguments may be IP addresses or hostnames, provided DNS name resolution is active. In order to reach this stage, a client must first have passed all of the standard connection tests in the control body.

8.1.3 deny

Type: slist

Allowed input range: (arbitrary string)

Synopsis: List of host names or IP addresses to deny access to file objects

Example:

Fill me in (`/home/mark/LapTop/Cfengine3/trunk/docs/bodypart_deny_example.texinfo`)

Notes:

Fill me in (`/home/mark/LapTop/Cfengine3/trunk/docs/bodypart_deny_notes.texinfo`)

8.1.4 maproot

Type: slist

Allowed input range: (arbitrary string)

Synopsis: List of host names or IP addresses to grant full read-privilege on the server

Example:

```
access:
```

```
"/home"
```

```
admit => { "backup_host.example.org" },
ifencrypted => "true",
maproot => "true";
```

Notes:

Normally users authenticated by the server are granted access only to files owned by them and no-one else. Even if the `cf-serverd` process runs with root privileges on the server side of a client-server connection, the client is not automatically granted access to download files owned by non-privileged users. If `maproot` is true then remote `root` users are granted access to all files.

A typical case where mapping is important is in making backups of many user files.

8.1.5 ifencrypted

Type: (menu option)

Allowed input range:

```

true
false
yes
no
on
off

```

Synopsis: true/false whether the current file access promise is conditional on the connection from the client being encrypted

Example:

`access:`

```

"/path/file"

admit    => { ".*.example.org" },
ifencrypted => "true";

```

Notes:

If this flag is true a client cannot access the file object unless its connection is encrypted.

8.2 roles promises in ‘server’

Roles promises are server-side decisions about which users are allowed to define soft-classes on the server’s system during remote invocation of `cf-agent`. This implements a form of Role Based Access Control (RBAC) for pre-assigned class-promise bindings. The user names cited must be attached to trusted public keys in order to be accepted.

```
roles:

  "regex"

    authorize => { "usernames", ... };
```

It is worth re-iterating here that it is not possible to send commands or modify promise definitions by remote access. At best users may try to send classes when using `cf-runagent` in order to activate sleeping promises. This mechanism limits their ability to do this.

```
bundle server access_rules()

{
roles:

  # Allow mark

  "Mark_.*" authorize => { "mark" };
}
```

In this example user `'mark'` is granted permission to remotely activate classes matching the regular expression when using the `cf-runagent` to activate cfengine. In this way one can implement a form of Role Based Access Control (RBAC), provided users do not have privileged access on the host directly.

8.2.1 authorize

Type: `slist`

Allowed input range: (arbitrary string)

Synopsis: List of public-key user names that are allowed to activate the promised class during remote agent activation

Example:

```
roles:

  ".*" authorize => { "mark", "marks_friend" };
```

Notes:

Part of Role Based Access Control (RBAC) in cfengine. The users listed in this section are granted access to set certain classes by using the remote `cf-runagent`. The user-names will refer to public key identities already trusted on the system.

9 Bundles of knowledge

```

bundle knowledge system
{
  topics:

  Troubleshooting::

  "Segmentation fault"
    association => a("is caused by","Bad memory reference","can cause");

  "Remote connection problem";
  "Web server not running";
  "Print server not running";
  "Bad memory reference";
}

```

Knowledge bundles describe topic maps, i.e. Topics, Associations and Occurrences (of topics in documents). This is for knowledge modelling and has no functional effect on a system.

9.1 topics promises in ‘knowledge’

Topic promises are part of the knowledge management engine. A topic is any string that refers to a concept or subject that we wish to include in a knowledge base. If a topic has a very long name, it is best to made the promiser object a short name and use the `comment` field to add the long explanation (e.g. unique acronym and full text).

```

topics:

"topic string"

  comment => "long name..",
  ...;

```

Topics form associative structures based entirely on an abstract space of natural language. Actually, this is only slightly more abstract than files, processes and commands etc. The main difference in knowledge management is that there are no corrective or maintenance operations associated with knowledge promises.

Class membership in knowledge management is subtly different from other parts of cfengine. If a topic lies in a certain class context, the topic uses it as a type-label. This is used for disambiguation of subject-area in searches rather than for disambiguation of rules between physical environments.

```

bundle knowledge example
{
  topics:

    "Distro"
      comment      => "Distribution of linux",
      association => a("is a packaging of","Linux","is packaged as a");
}

```

Topics are basically identifiers, where the comment field here is a long form of the subject string. Associations form semantic links between topics. Topics can appear multiple times in order to form multiple associations.

9.1.1 association (compound body)

Type: (ext body)

`'forward_relationship'`

Type: string

Allowed input range: (arbitrary string)

Synopsis: Name of forward association between promiser topic and associates

Example:

```

body association example
{
  forward_relation => "is bigger than";
}

```

Notes:

`'backward_relationship'`

Type: string

Allowed input range: (arbitrary string)

Synopsis: Name of backward/inverse association from associates to promiser topic

Example:

```
body association example
{
# ..
backward_relationship => "is less than";
}
```

Notes:

Denotes the inverse name which is used to ‘moralizing’ the association graph.

‘associates’

Type: slist

Allowed input range: (arbitrary string)

Synopsis: List of associated topics by this forward relationship

Example:

```
body association example(literal,scalar,list)

{
#...
associates => { "literal", $(scalar), @(list)};
}
```

Notes:

An element of an association which is a list of topics to which the current topic is associated.

9.1.2 comment

Type: string

Allowed input range: (arbitrary string)

Synopsis: Retained comment about this promise’s real intention

Example:

```
comment => "This comment follows the data for reference ...",
```

Notes:

Comments written in code follow the program, they are not merely discarded. They appear in reports and error messages.

9.2 occurrences promises in 'knowledge'

Occurrences are documents or information resources that discuss topics. An occurrence promise asserts that a particular document or text resource in fact represents information about one or more topics. This is used to construct references to actual information in a topic map.

```

occurrences:
  topic_name::
    "URL reference or literal string"
    represents => { "sub-topic disambiguator", ... },
    representation => "literal or url";

```

Mark_Burgess::

```

"http://www.iu.hio.no/~mark"
  represents => { "Home Page" };

```

lvalue::

```

"A variable identifier, i.e. the left hand side of an '=' association. The promiser in a variable
  represents => { "Definitions" },
  representation => "literal";

```

Editing_Files::

```

"http://www.cfengine.org/confdir/customizepasswd.html"
  represents => { "Setting up users" };

```

Occurrences are pointers to information about topics. This might be a literal text string or a URL reference to an external document.

9.2.1 represents

Type: slist

Allowed input range: (arbitrary string)

Synopsis: List of subtopics that disambiguate the context of this reference

Example:

occurrences:

```
Promise_Theory::
```

```
"A theory of autonomous actors that offer certainty through promises"
```

```
  represents    => { "Definitions" },
  representation => "literal";
```

Notes:

The sub-topic represented by the document reference in a knowledge base.

9.2.2 representation

Type: (menu option)

Allowed input range:

```
  literal
  url
  db
  file
  web
```

Synopsis: How to interpret the promiser string e.g. actual data or reference to data

Example:

occurrences:

```
Promise_Theory::
```

```
"A theory of autonomous actors that offer certainty through promises"
```

```
  represents    => { "Definitions" },
  representation => "literal";
```

Notes:

The form of knowledge representation in a topic map occurrence reference.

10 Special functions

10.1 Function accessedbefore

Synopsis: `accessedbefore(2 args)` returns type class

Define class if `arg1` was accessed before `arg2` (`atime`)

Example:

```
body common control

{
bundlesequence => { "example" };
}

#####

bundle agent example

{
classes:

  "do_it" and => { accessedbefore("/tmp/earlier","/tmp/later"), "linux" };

reports:

  do_it::

    "The secret changes have been accessed after the reference time";

}
```

Notes:

The function accesses the `atime` fields of a file and makes a comparison.

```
touch /tmp/reference
touch /tmp/secretfile

/usr/local/sbin/cf-agent -f ./unit_accessed_before.cf -K
R: The secret changes have been accessed after the reference time
```

10.2 Function accumulated

Synopsis: `accumulated(6 args)` returns type int

Convert an accumulated amount of time into a system representation

Example:

```

bundle agent testbundle

{
processes:

    ".*"

    process_count    => anyprocs,
    process_select  => proc_finder;

reports:

    any_procs::

        "Found processes in range";
}

#####

body process_select proc_finder

{
ttime_range => irange(accumulated(0,0,0,0,2,0),accumulated(0,0,0,0,20,0));
process_result => "ttime";
}

#####

body process_count anyprocs

{
match_range => "0,0";
out_of_range_define => { "any_procs" };
}

```

Notes:

In the example we look for processes that have accumulated between 2 and 20 minutes of total run time.

ARGUMENTS:

'Years' The year, e.g. 2009
 'Month' The Month, 1-12
 'Day' The day 1-31
 'Hours' The hour 0-23
 'Minutes' The minutes 0-59
 'Seconds' The number of seconds 0-59

10.3 Function ago

Synopsis: ago(6 args) returns type int

Convert a time relative to now to an integer system representation

Example:

```

bundle agent testbundle

{
processes:

  "."

  process_count  => anyprocs,
  process_select => proc_finder;

reports:

  any_procs::

    "Found processes out of range";
}

#####

body process_select proc_finder

{
stime_range => irange(ago(0,0,0,5,30,0),ago(0,0,0,0,20,0));
process_result => "stime";
}

#####

body process_count anyprocs
  
```

```
{
match_range => "0,0";
out_of_range_define => { "any_procs" };
}
```

Notes:

The ago function measures time relative to now.

ARGUMENTS:

'Years' The year, e.g. 2009
'Months' The Month, 1-12
'Days' The day 1-31
'Hours' The hour 0-23
'Minutes' The minutes 0-59
'Seconds' The number of seconds 0-59

10.4 Function canonify

Synopsis: canonify(1 args) returns type string
Convert an arbitrary string into a legal class name

Example:

commands:

```
"/var/cfengine/bin/${component}"

ifvarclass => canonify("start_${component}");
```

Notes:

This is for use in turning arbitrary text into class data.

10.5 Function changedbefore

Synopsis: changedbefore(2 args) returns type class
Define class if arg1 was changed before arg2 (ctime)

Example:

```

body common control

{
bundlesequence => { "example" };
}

#####

bundle agent example

{
classes:

  "do_it" and => { changedbefore("/tmp/earlier","/tmp/later"), "linux" };

reports:

  do_it::

    "The derived file needs updating";

}

Notes:

```

Change times include both file permissions and file contents. Comparisons like this are normally used for updating files (like the ‘make’ command).

10.6 Function classmatch

Synopsis: classmatch(1 args) returns type class

Define class if the regular expression matches any currently defined class

Example:

```

body common control

{
bundlesequence => { "example" };
}

#####

bundle agent example

```

```

{
classes:

  "do_it" and => { classmatch(".*_cfengine_com"), "linux" };

reports:

  do_it::

    "Host matches pattern";
}

```

Notes:

The regular expression is matched against the current list of defined classes.

10.7 Function execresult

Synopsis: execresult(2 args) returns type string

Execute named command and assign output to variable

Example:

body common control

```

{
bundlesequence => { "example" };
}

```

#####

bundle agent example

```

{
vars:

  "my_result" string => execresult("/bin/ls /tmp","noshell");

reports:

  linux::

    "Variable is $(my_result)";
}

```



```
}
```

Notes:

The second argument decides whether a shell will be used to encapsulate the command. This is necessary in order to combine commands with pipes etc, but remember that each command requires a new process that reads in files beyond cfengine's control. Thus using a shell is both a performance hog and a potential security issue.

10.8 Function fileexists

Synopsis: fileexists(1 args) returns type class

Define class if the named file can be accessed

Example:

```
body common control

{
bundlesequence => { "example" };
}

#####

bundle agent example

{
classes:

    "exists" expression => fileexists("/etc/passwd");

reports:

    exists::

        "File exists";
}

```

Notes:

The user must have access permissions to the file for this to work faithfully.

10.9 Function getindices

Synopsis: getindices(1 args) returns type slist

Get a list of keys to the array whose id is the argument and assign to variable

Example:

```
body common control

{
  any::

    bundlesequence => { "testsetvar" };
}

#####

bundle agent testsetvar

{
  vars:

    "v[index_1]" string => "value_1";
    "v[index_2]" string => "value_2";

    "parameter_name" slist => getindices("v");

  reports:

    Yr2008::

      "Found index: $(parameter_name)";

}
```

Notes:

Make sure you specify the correct scope when supplying the name of the variable.

10.10 Function getgid

Synopsis: getgid(1 args) returns type int

Return the integer group id of the named group on this host

Example:

```

body common control

{
bundlesequence => { "example" };
}

#####

bundle agent example

{
vars:

    "gid" int => getgid("users");

reports:

    Yr2008::

        "Users gid is $(gid)";

}

```

Notes:

If the named group does not exist, the variable will not be defined.

10.11 Function getuid

Synopsis: getuid(1 args) returns type int

Return the integer user id of the named user on this host

Example:

```

body common control

{
bundlesequence => { "example" };
}

#####

```

```

bundle agent example

{
vars:

    "uid" int => getuid("mark");

reports:

    Yr2008::

        "Users gid is ${uid}";

}

```

Notes:

If the named user is not registered the variable will not be defined.

10.12 Function groupexists

Synopsis: groupexists(1 args) returns type class
 Define class if group or numerical id exists on this host

Example:

```

body common control

{
bundlesequence => { "example" };
}

#####

bundle agent example

{
classes:

    "gname" expression => groupexists("users");
    "gid"    expression => groupexists("100");

reports:

```

```

gname::
    "Group exists by name";

gid::
    "Group exists by id";

}

```

Notes:

The group may be specified by name or number.

10.13 Function hash

Synopsis: hash(2 args) returns type string

Return the hash of arg1, type arg2 and assign to a variable

Example:

```

body common control

{
bundlesequence => { "example" };
}

#####

bundle agent example

{
vars:

    "md5" string => hash("Cfengine is not cryptic","md5");

reports:

    Yr2008::

        "Hashed to: $(md5)";

}

```

Notes:

Hash functions are extremely sensitive to input. You should not expect to get the same answer from this function as you would from every other tool, since it depends on how whitespace and end of file characters are handled.

10.14 Function hostrange

Synopsis: hostrange(2 args) returns type class

Define class if the current host lies in the range of enumerated hostnames specified

Example:

```
body common control

{
bundlesequence => { "example" };
}

#####

bundle agent example

{
classes:

  "compute_nodes" expression => hostrange("cpu-", "01-32");

reports:

  compute_nodes::

    "No computer is a cluster";

}
```

Notes:

This is a pattern matching function for non-regular (enumerated) expressions.

10.15 Function iprange

Synopsis: iprange(1 args) returns type class

Define class if the current host lies in the range of IP addresses specified

Example:

```
body common control

{
bundlesequence => { "example" };
}

#####

bundle agent example

{
classes:

    "adhoc_group_1" expression => iprange("128.39.89.10-15");
    "adhoc_group_2" expression => iprange("128.39.74.1/23");

reports:

    adhoc_group_1::

        "Some numerology";

    adhoc_group_2::

        "The masked warriors";
}

```

Notes:

Pattern matching based on IP addresses.

10.16 Function `irange`

Synopsis: `irange(2 args)` returns type `irange [int,int]`

Define a range of integer values for cfengine internal use

Example:

Notes:

Not currently used.

10.17 Function isdir

Synopsis: isdir(1 args) returns type class

Define class if the named object is a directory

Example:

```
body common control

{
bundlesequence => { "example" };
}

#####

bundle agent example

{
classes:

  "isdir" expression => isdir("/etc");

reports:

  isdir::

    "Directory exists..";
}

```

Notes:

The cfengine process must have access to the object concerned in order for this to work.

10.18 Function isgreaterthan

Synopsis: isgreaterthan(2 args) returns type class

Define class if arg1 is numerically greater than arg2, else compare strings like strcmp

Example:

```

body common control

{
bundlesequence => { "test" };
}

#####

bundle agent test

{
classes:

  "ok" expression => isgreaterthan("1","0");

reports:

  ok::

    "Assertion is true";

  !ok::

    "Assertion is false";

}

```

Notes:

The comparison is made numerically if possible. If the values are strings, the result is identical to that of comparing with `strcmp()`.

10.19 Function islessthan

Synopsis: `islessthan(2 args)` returns type class

Define class if `arg1` is numerically less than `arg2`, else compare strings like NOT `strcmp`

Example:

```

body common control

```

```
{
bundlesequence => { "test" };
}
```

```
#####
```

```
bundle agent test
```

```
{
classes:

    "ok" expression => islessthan("0","1");

reports:

    ok::

        "Assertion is true";

    !ok::

        "Assertion is false";

}
```

Notes:

The complement of `isgreaterthan`. The comparison is made numerically if possible. If the values are strings, the result is identical to that of comparing with `'strcmp()'`.

10.20 Function `islink`

Synopsis: `islink(1 args)` returns type class
 Define class if the named object is a symbolic link

Example:

```
body common control

{
bundlesequence => { "example" };
}
```

```
#####
```

```

bundle agent example

{
classes:

  "isdir" expression => islink("/tmp/link");

reports:

  isdir::

    "Directory exists..";
}

```

Notes:

The link node must both exist and be a symbolic link. Hard links cannot be detected using this function. A hard link is a regular file or directory.

10.21 Function isnewerthan

Synopsis: isnewerthan(2 args) returns type class

Define class if arg1 is newer (modified later) than arg2 (mtime)

Example:

```

body common control

{
bundlesequence => { "example" };
}

#####

bundle agent example

{
classes:

  "do_it" and => { isnewerthan("/tmp/later", "/tmp/earlier"), "linux" };

reports:

  do_it::

```

```

    "The derived file needs updating";
}

```

Notes:

This function compares the modification time of the file, referring to changes of content only.

10.22 Function isplain

Synopsis: isplain(1 args) returns type class

Define class if the named object is a plain/regular file

Example:

```

body common control

{
bundlesequence => { "example" };
}

#####

bundle agent example

{
classes:

    "isplain" expression => isplain("/etc/passwd");

reports:

    isplain::

        "File exists..";

}

```

Notes:

10.23 Function isvariable

Synopsis: `isvariable(1 args)` returns type class

Define class if the named variable is defined

Example:

```
body common control

{
bundlesequence => { "example" };
}

#####

bundle agent example

{
vars:

    "bla" string => "xyz..";

classes:

    "exists" expression => isvariable("bla");

reports:

    exists::

        "Variable exists: \"$(bla)\"..";

}
```

Notes:

The variable need only exist. This says nothing about its value. Use `regcmp` to check variable values.

10.24 Function now

Synopsis: `now(0 args)` returns type int

Convert the current time into system representation

Example:

```
body file_select zero_age
{
mtime      => irange(ago(1,0,0,0,0,0),now);
file_result => "mtime";
}
```

Notes:**10.25 Function on**

Synopsis: on(6 args) returns type int

Convert an exact date/time to an integer system representation

Example:

```
body file_select zero_age
{
mtime      => irange(on(2000,1,1,0,0,0),now);
file_result => "mtime";
}
```

Notes:

An absolute date. Note that in process matching dates could be wrong by an hour depending on Daylight Savings Time / Summer Time. This is a known bug to be fixed.

ARGUMENTS:

'Years'	The year, e.g. 2009
'Month'	The Month, 1-12
'Day'	The day 1-31
'Hours'	The hour 0-23
'Minutes'	The minutes 0-59
'Seconds'	The number of seconds 0-59

10.26 Function randomint

Synopsis: randomint(2 args) returns type int

Generate a random integer between the given limits

Example:

vars:

```
"ran"    int => randomint(4,88);
```

Notes:

The limits must be integer values and the resulting numbers are based on the entropy of the md5 algorithm.

10.27 Function readfile

Synopsis: readfile(2 args) returns type string

Read max number of bytes from named file and assign to variable

Example:

vars:

```
"xxx"

string => readfile( "/home/mark/tmp/testfile" , "33" );
```

Notes:

The file (fragment) is read into a single scalar variable.

10.28 Function readintarray

Synopsis: readintarray(6 args) returns type int

Read an array of integers from a file and assign the dimension to a variable

Example:

vars:

```
"dim_array"

int => readintarray("array_name", "/tmp/array", "#[^\n]*", ":", 10, 4000);
```

Notes:

Reads a two dimensional array from a file. One dimension is separated by the character specified in the argument, the other by the the lines in the file. The first field of the lines names the first array argument.

```
1: 5:7:21:13
2:19:8:14:14
3:45:1:78:22
4:64:2:98:99
```

Results in

```
array_name[1][0] 1
array_name[1][1] 5
array_name[1][2] 7
array_name[1][3] 21
array_name[1][4] 13
array_name[2][0] 2
array_name[2][1] 19
array_name[2][2] 8
array_name[2][3] 14
array_name[2][4] 14
array_name[3][0] 3
array_name[3][1] 45
array_name[3][2] 1
array_name[3][3] 78
array_name[3][4] 22
array_name[4][0] 4
array_name[4][1] 64
array_name[4][2] 2
array_name[4][3] 98
array_name[4][4] 99
```

10.29 Function readintlist

Synopsis: readintlist(5 args) returns type ilet

Read and assign a list variable from a file of separated ints

Example:

```
body common control
```

```
{
bundlesequence => { "example" };
}
```

```
#####
```



```
bundle agent example
```

```
{
vars:

    "mylist" ilist => { readintlist("/tmp/listofint", "#.*", "[\n]", 10, 400) };

reports:

    Yr2008::

        "List entry: ${mylist}";
}
```

ARGUMENTS:

‘filename’ The name of a text file containing text to be split up as a list.

‘comment’ A regex pattern which is to be ignored in the file

‘split’ A regex pattern which is to be used to split up the file into items

‘maxent’ The maximum number of list items to read from the file

‘maxsize’ The maximum number of bytes to read from the file

Notes:

10.30 Function readrealarray

Synopsis: readrealarray(6 args) returns type int

Read an array of real numbers from a file and assign the dimension to a variable

Example:

```
vars:

    "dim_array"

    int => readrealarray("array_name", "/tmp/array", "#[^\n]*", ":", 10, 4000);
```

Notes:

See the notes for readintarray.

10.31 Function readreallist

Synopsis: readreallist(5 args) returns type rlist

Read and assign a list variable from a file of separated real numbers

Example:

```
body common control
```

```
{
bundlesequence => { "example" };
}
```

```
#####
```

```
bundle agent example
```

```
{
vars:

    "mylist" ilist => { readreallist("/tmp/listofreal", "#.*", "[\n]", 10, 400) };

reports:

    Yr2008::

        "List entry: ${mylist}";

}
```

ARGUMENTS:

- 'filename' The name of a text file containing text to be split up as a list.
- 'comment' A regex pattern which is to be ignored in the file
- 'split' A regex pattern which is to be used to split up the file into items
- 'maxent' The maximum number of list items to read from the file
- 'maxsize' The maximum number of bytes to read from the file

Notes:

10.32 Function readstringarray

Synopsis: readstringarray(6 args) returns type int

Read an array of strings from a file and assign the dimension to a variable

Example:

vars:

```
"dim_array"

int => readstringarray("array_name", "/tmp/array", "#[^\n]*", ":", 10, 4000);
```

Notes:

Reads a two dimensional array from a file. One dimension is separated by the character specified in the argument, the other by the the lines in the file. The first field of the lines names the first array argument.

```
at:x:25:25:Batch jobs daemon:/var/spool/atjobs:/bin/bash
avahi:x:103:105:User for Avahi:/var/run/avahi-daemon:/bin/false
beagleindex:x:104:106:User for Beagle indexing:/var/cache/beagle:/bin/bash
bin:x:1:1:bin:/bin:/bin/bash
daemon:x:2:2:Daemon:/sbin:/bin/bash
```

Results in a systematically indexed map of the file. Some samples are show below to illustrate the pattern.

```
...
array_name[daemon][0] daemon
array_name[daemon][1] x
array_name[daemon][2] 2
array_name[daemon][3] 2
array_name[daemon][4] Daemon
array_name[daemon][5] /sbin
array_name[daemon][6] /bin/bash
...
array_name[at][3] 25
array_name[at][4] Batch jobs daemon
array_name[at][5] /var/spool/atjobs
array_name[at][6] /bin/bash
...
array_name[games][3] 100
array_name[games][4] Games account
array_name[games][5] /var/games
array_name[games][6] /bin/bash
...
```

10.33 Function readstringlist

Synopsis: readstringlist(5 args) returns type slist

Read and assign a list variable from a file of separated strings

Example:

```

body common control

{
bundlesequence => { "example" };
}

#####

bundle agent example

{
vars:

    "mylist" ilist => { readstringlist("/tmp/listofint", "#.*", "[\n]", 10, 400) };

reports:

    Yr2008::

        "List entry: $(mylist)";

}

```

ARGUMENTS:

‘filename’ The name of a text file containing text to be split up as a list.

‘comment’ A regex pattern which is to be ignored in the file

‘split’ A regex pattern which is to be used to split up the file into items

‘maxent’ The maximum number of list items to read from the file

‘maxsize’ The maximum number of bytes to read from the file

Notes:

The following example file would be split into a list of the first ten Green letters.

```

alpha
beta
gamma # This is a comment
delta
epsilon
zeta
eta
theta

```

```

iota
kappa
lambda
mu
nu
etc

```

10.34 Function readtcp

Synopsis: readtcp(4 args) returns type string

Connect to tcp port, send string and assign result to variable

Example:

```

body common control

{
bundlesequence => { "example" };
}

#####

bundle agent example

{
vars:

    "my80" string => readtcp("www.cfengine.com","80","GET /index.html",400);

reports:

    Yr2008::

        "Server returned: $(my80)";

}

```

‘hostnameip’ The host name or IP address of a tcp socket.

‘port’ The port number to connect to.

‘sendstring’ A string to send to the TCP port to illicit a response

‘maxbytes’ The maximum number of bytes to read in response.

Notes:

If the send string is empty, no data are sent or received from the socket. Then the function only tests whether the TCP port is alive and returns an empty variable.

Note that on some systems the timeout mechanism does not seem to successfully interrupt the waiting system calls so this might hang if you send a query string that is incorrect. This should not happen, but the cause has yet to be diagnosed.

10.35 Function regarray

Synopsis: `regarray(2 args)` returns type class

Define class if `arg1` matches any item in the associative array with `id=arg2`

Example:

```
body common control

{
bundlesequence => { "testbundle" };
}

#####

bundle agent testbundle
{
vars:

    "myarray[0]" string => "bla1";
    "myarray[1]" string => "bla2";
    "myarray[3]" string => "bla";
    "myarray"    string => "345";
    "not"        string => "345";

classes:

    "ok" expression => regarray("myarray","b.*2");

reports:

    ok::

        "Found in list";

    !ok::
```

```

    "Not found in list";
}

```

Notes:

Tests whether an associative array contains elements matching a certain regular expression. The result is a class.

ARGUMENTS:

`'array_name'` The name of the array, with no `'$()'` surrounding, etc.
`'regex'` A regular expression to match the content.

10.36 Function regcmp

Synopsis: `regcmp(2 args)` returns type class

Define class if `arg1` is a regular expression matching `arg2`

Example:

```

bundle agent subtest(user)

{
classes:

    "invalid" not => regcmp("[a-z][a-z][a-z][a-z]", "$(user)");

reports:

!invalid::

    "User name $(user) is valid at 4 letters";

invalid::

    "User name $(user) is invalid";
}

```

Notes:

Compares a string to a regular expression.

ARGUMENTS:

`'regex'` A regular expression to match the test data.

`'string'` Test data for the regular expression.

10.37 Function `reglist`

Synopsis: `reglist(2 args)` returns type class

Define class if `arg1` matches any item in the list with `id=arg2`

Example:

vars:

```
"nameservers" slist => {
    "128.39.89.10",
    "128.39.74.16",
    "192.168.1.103"
};
```

classes:

```
"am_name_server" expression => reglist("@(nameservers)", "${sys.ipv4[eth0]}");
```

Notes:

Matches a list of test strings to a regular expression.

ARGUMENTS:

`'list'` The list of strings.

`'regex'` The scalar regular expression string.

10.38 Function `returnszero`

Synopsis: `returnszero(2 args)` returns type class

Define class if named shell command has exit status zero

Example:

body common control

```
{
bundlesequence => { "example" };
}
```

```
#####
```


bundle agent example

```
{
classes:

    "my_result" expression => returnszero("/usr/local/bin/mycommand","noshell");

reports:

    !my_result::

        "Command failed";

}
```

Notes:

This is the complement of `execresult`, but it returns a class result rather than the output of the command.

10.39 Function `rrange`

Synopsis: `rrange(2 args)` returns type `range [real,real]`
 Define a range of real numbers for cfengine internal use

Example:

?

Notes:

This is not yet used.

10.40 Function `selectservers`

Synopsis: `selectservers(6 args)` returns type `int`
 Select tcp servers which respond correctly to a query and return their number, set array of names

Example:

body common control

```

{
bundlesequence => { "test" };
}

#####

bundle agent test

{
vars:

    "hosts" slist => { "slogans.iu.hio.no", "eternity.iu.hio.no", "nexus.iu.hio.no" };

    "up_servers" int => selectservers("@(hosts)","80","","","100","alive_servers");

classes:

    "someone_alive" expression => isgreaterthan("${up_servers}","0");

    "i_am_a_server" expression => regarray("up_servers","$(host)|$(fqhost)");

reports:

    someone_alive::

        "Number of active servers ${up_servers}";
        "First server ${alive_servers[0]} fails over to ${alive_servers[1]}";

}

```

Notes:

This function selects all the TCP ports that are active and functioning from an ordered list and builds an array of their names. This allows us to select a current list of failover alternatives that are pretested.

'hostlist' A list of host names or IP addresses to attempt to connect to.

'port' The port number for the service.

'sendstr' An optional string to send to the server to illicit a response.

'regex_on_reply'

If a string is sent, this regex must match the resulting reply.

'maxbytesread_reply'

The maximum number of bytes to read as the server's reply.

`'array_name'`

The name of the array to build containing the names of hosts that pass the above tests.
The array is ordered `array_name[0], .. etc.`

10.41 Function strcmp

Synopsis: `strcmp(2 args)` returns type class

Define class if the two strings match exactly

Example:

```
body common control

{
bundlesequence => { "example" };
}

#####

bundle agent example

{
classes:

  "same" expression => strcmp("test","test");

reports:

  same::

    "Strings are equal";

  !same::

    "Strings are not equal";
}
```

Notes:

10.42 Function usemodule

Synopsis: `usemodule(2 args)` returns type class

Execute cfengine module script and set class if successful

Example:

```

body common control
{
  any::

    bundlesequence => {
                        test
                    };
}

#####

bundle agent test

{
classes:

  # returns $(user)

  "done" expression => usemodule("getusers","");

commands:

  "/bin/echo promiser text" args => "test $(user)";
}

```

Notes:

Modules must reside in 'WORKDIR/modules' but no longer require a special naming convention.

ARGUMENTS:**'Module name'**

The name of the module without its leading path, since it is assumed to be in the registered modules directory.

'Argument string'

Any command link arguments to pass to the module.

10.43 Function userexists

Synopsis: userexists(1 args) returns type class

Define class if user name or numerical id exists on this host

Example:

```
body common control

{
bundlesequence => { "example" };
}

#####

bundle agent example

{
classes:

  "ok" expression => userexists("root");

reports:

  ok::

    "Root exists";

  !ok::

    "Root does not exist";
}
```

Notes:

Checks whether the user is in the password database for the current host. The argument must be a user name or user id.

11 Special Variables

11.1 Variable context const

Cfengine defines a number of variables for embedding unprintable values or values with special meanings in strings.

11.1.1 Variable const.dollar

reports:

some::

```
"The value of $(const.dollar)(const.dollar) is $(const.dollar)";
```

```
"But the value of \$(dollar) is \$(dollar)";
```

11.1.2 Variable const.endl

reports:

```
"A newline with either $(const.endl) or with $(const.n) is ok";
```

11.1.3 Variable const.n

reports:

```
"A newline with either $(const.n) or with $(const.endl) is ok";
```

11.1.4 Variable const.r

reports:

```
"A carriage return character is $(const.r)";
```

11.2 Variable context sys

System variables are derived from cfengine's automated discovery of system values. They are provided as variables in order to make automatically adaptive rules for configuration, e.g.

```
files:
```

```
any::
```

```
"$(sys.resolve)"

    create      => "true",
    edit_line   => doresolve("@(this.list1)","@(this.list2)"),
    edit_defaults => reconstruct;
```

The above rule requires no class specification because the variable itself is class-specific.

11.2.1 Variable sys.arch

The variable gives the kernel's short architecture description.

```
# arch = x86_64
```

11.2.2 Variable sys.cdate

The date of the system in canonical form, i.e. in the form of a class.

```
# cdate = Sun_Dec__7_10_39_53_2008_
```

11.2.3 Variable sys.class

This variable contains the name of the hard-class category for this host, i.e. its top level operating system type classification.

```
# class = linux
```


11.2.4 Variable `sys.date`

The date of the system as a text string.

```
# date = Sun Dec 7 10:39:53 2008
```

11.2.5 Variable `sys.domain`

The domain name as divined by `cfengine`. If the DNS is in use, it could be possible to derive the domain name from its DNS registration, but in general there is no way to discover this value automatically. The `common control` body permits the ultimate specification of this value.

```
# domain = example.org
```

11.2.6 Variable `sys.fqhost`

The fully qualified name of the host. In order to compute this value properly, the domain name must be defined.

```
# fqhost = host.example.org
```

11.2.7 Variable `sys.fstab`

The location of the system filesystem (mount) table.

```
# fstab = /etc/fstab
```

11.2.8 Variable `sys.host`

The name of the current host, according to the kernel. It is undefined whether this is qualified or unqualified with a domain name.

```
# host = myhost
```

11.2.9 Variable `sys.interface`

The assumed (default) name of the main system interface on this host.

```
# interface = eth0
```

11.2.10 Variable `sys.ipv4`

All four octets of the IPv4 address of the system interface named as the associative array index, e.g. `'${ipv4_1[${interface}]}'`.

```
# The IP on the default interface
#     ipv4 = 192.168.1.101

# The octets of all interfaces as an associative array
# ipv4_1[1e0] = 192
# ipv4_2[1e0] = 192.168
# ipv4_3[1e0] = 192.168.1
#   ipv4[1e0] = 192.168.1.101
```

11.2.11 Variable `sys.ipv4[eth1]`

The IPv4 address of the default system interface.

```
# The IP on the default interface
#     ipv4 = 192.168.1.101

# The octets of all interfaces as an associative array
# ipv4_1[1e0] = 192
# ipv4_2[1e0] = 192.168
# ipv4_3[1e0] = 192.168.1
#   ipv4[1e0] = 192.168.1.101
```

11.2.12 Variable `sys.ipv4_1[eth1]`

The first octet of the IPv4 address of the system interface named as the associative array index, e.g. `'${ipv4_1[${interface}]}'`.

```
# The IP on the default interface
#     ipv4 = 192.168.1.101

# The octets of all interfaces as an associative array
# ipv4_1[1e0] = 192
# ipv4_2[1e0] = 192.168
# ipv4_3[1e0] = 192.168.1
#   ipv4[1e0] = 192.168.1.101
```

11.2.13 Variable `sys.ipv4_2[eth1]`

The first two octets of the IPv4 address of the system interface named as the associative array index, e.g. `'${ipv4_1[${interface}]}'`.

```
# The IP on the default interface
#     ipv4 = 192.168.1.101

# The octets of all interfaces as an associative array
# ipv4_1[1e0] = 192
# ipv4_2[1e0] = 192.168
# ipv4_3[1e0] = 192.168.1
#   ipv4[1e0] = 192.168.1.101
```

11.2.14 Variable `sys.ipv4_3[eth1]`

The first three octets of the IPv4 address of the system interface named as the associative array index, e.g. `'${ipv4_1[${interface}]}'`.

```
# The IP on the default interface
#     ipv4 = 192.168.1.101

# The octets of all interfaces as an associative array
# ipv4_1[1e0] = 192
# ipv4_2[1e0] = 192.168
# ipv4_3[1e0] = 192.168.1
#   ipv4[1e0] = 192.168.1.101
```

11.2.15 Variable `sys.long_arch`

The long architecture name for this system kernel. This name is sometimes quite unwieldy but can be useful for logging purposes.

```
# long_arch = linux_x86_64_2_6_22_19_0_1_default__1_SMP_2008_10_14_22_17_43__0200
```

11.2.16 Variable `sys.maildir`

The name of the system email spool directory.

```
# maildir = /var/spool/mail
```

11.2.17 Variable `sys.os`

The name of the operating system according to the kernel.

```
# os = linux
```

11.2.18 Variable `sys.ostype`

Another name for the operating system.

```
# ostype = linux_x86_64
```

11.2.19 Variable `sys.release`

The kernel release of the operating system.

```
# release = 2.6.22.19-0.1-default
```

11.2.20 Variable `sys.resolv`

The location of the system resolver file.

```
# resolv = /etc/resolv.conf
```

11.2.21 Variable `sys.uqhost`

The unqualified name of the current host. See also `sys.fqhost`.

```
# uqhost = myhost
```

11.2.22 Variable `sys.workdir`

The location of the cfengine work directory and cache. For the system privileged user this is normally:

```
# workdir = /var/cfengine
```

For non-privileged users it is in the user's home directory:

```
# workdir = /home/user/.cfagent
```


Table of Contents

1	Cfengine 3.0.0b5 Basics	1
1.1	Why change the language?	2
1.2	Testing as a non-privileged user	3
1.3	The bare necessities of a cfengine 3	3
1.4	Familiarizing yourself	4
2	A simple crash course in concepts	9
2.1	Rules are promises	9
2.2	Containers	10
2.3	Types in cfengine 3	11
3	How to run cfengine 3 examples	13
4	A complete configuration	15
4.1	What should a failsafe file contain?	19
4.2	Recovery from errors in the configuration	19
4.3	Recovery from errors in the software	21
5	Control promises	23
5.1	common control promises	23
5.1.1	bundlesequence	23
5.1.2	inputs	24
5.1.3	version	24
5.1.4	lastseenexpireafter	25
5.1.5	output_prefix	25
5.1.6	domain	25
5.2	agent control promises	26
5.2.1	maxconnections	26
5.2.2	abortclasses	27
5.2.3	abortbundleclasses	27
5.2.4	addclasses	29
5.2.5	agentaccess	29
5.2.6	agentfacility	30
5.2.7	auditing	30
5.2.8	binarypaddingchar	31
5.2.9	bindtointerface	31
5.2.10	hashpurge	32
5.2.11	hashupdates	32
5.2.12	childlibpath	33
5.2.13	defaultcopytype	33
5.2.14	dryrun	34
5.2.15	editbinaryfilesize	34

5.2.16	editfilesize	35
5.2.17	exclamation	35
5.2.18	expireafter	36
5.2.19	files_single_copy	36
5.2.20	files_auto_define	37
5.2.21	fullencryption	37
5.2.22	hostnamekeys	38
5.2.23	ifelapsed	38
5.2.24	inform	39
5.2.25	lastseen	39
5.2.26	lastseenexpireafter	40
5.2.27	mountfilesystems	40
5.2.28	nonalphanumfiles	41
5.2.29	repchar	42
5.2.30	default_repository	42
5.2.31	secureinput	42
5.2.32	sensiblecount	43
5.2.33	sensiblesize	43
5.2.34	skipidentify	44
5.2.35	suspiciousnames	44
5.2.36	syslog	45
5.2.37	timezone	45
5.2.38	default_timeout	45
5.2.39	verbose	46
5.3	server control promises	46
5.3.1	cfruncommand	47
5.3.2	maxconnections	47
5.3.3	denybadclocks	48
5.3.4	allowconnects	48
5.3.5	denyconnects	49
5.3.6	allowallconnects	49
5.3.7	trustkeysfrom	50
5.3.8	allowusers	50
5.3.9	dynamicaddresses	51
5.3.10	skipverify	51
5.3.11	logallconnections	51
5.3.12	logencryptedtransfers	52
5.3.13	hostnamekeys	53
5.3.14	auditing	53
5.3.15	bindtointerface	54
5.3.16	serverfacility	54
5.4	monitor control promises	55
5.4.1	forgetrate	55
5.4.2	monitorfacility	56
5.4.3	histograms	56
5.4.4	tcpdump	57
5.4.5	tcpdumpcommand	57
5.5	runagent control promises	58

5.5.1	hosts	58
5.5.2	port	59
5.5.3	force_ipv4	59
5.5.4	trustkey	60
5.5.5	encrypt	60
5.5.6	background_children	61
5.5.7	max_children	62
5.5.8	output_to_file	62
5.6	executor control promises	63
5.6.1	splaytime	63
5.6.2	mailfrom	63
5.6.3	mailto	64
5.6.4	smtpserver	64
5.6.5	mailmaxlines	65
5.6.6	schedule	65
5.6.7	executorfacility	65
5.6.8	execcommand	66
5.7	knowledge control promises	67
5.7.1	id_prefix	67
5.7.2	build_directory	68
5.7.3	sql_type	68
5.7.4	sql_database	68
5.7.5	sql_owner	69
5.7.6	sql_passwd	69
5.7.7	sql_server	70
5.7.8	query_output	70
5.7.9	query_engine	71
5.7.10	style_sheet	71
5.7.11	html_banner	71
5.7.12	graph_output	72
5.7.13	graph_directory	72
5.7.14	generate_manual	73
5.7.15	manual_source_directory	73
5.8	reporter control promises	74
5.8.1	reports	74
5.8.2	report_output	75
5.8.3	build_directory	75
5.8.4	auto_scaling	76
5.8.5	error_bars	76
5.8.6	time_stamps	77
5.8.7	query_engine	78
5.8.8	style_sheet	78
5.8.9	html_banner	78

6	Bundles of common	81
6.1	vars promises	81
6.1.1	string	81
6.1.2	int	82
6.1.3	real	82
6.1.4	slist	83
6.1.5	ilist	83
6.1.6	rlist	84
6.1.7	policy	84
6.2	classes promises	85
6.2.1	or	85
6.2.2	and	85
6.2.3	xor	86
6.2.4	dist	86
6.2.5	expression	87
6.2.6	not	87
6.3	reports promises	87
6.3.1	lastseen	87
6.3.2	intermittency	88
6.3.3	showstate	88
6.3.4	printfile (compound body)	91
6.4	* promises in 'agent'	92
6.4.1	action (compound body)	92
6.4.2	classes (compound body)	97
6.4.3	ifvarclass	101
6.4.4	comment	102
6.5	* promises in 'edit_line'	102
6.5.1	select_region (compound body)	103
7	Bundles of agent	105
7.1	commands promises in 'agent'	105
7.1.1	args	106
7.1.2	contain (compound body)	107
7.1.3	module	111
7.2	files promises in 'agent'	113
7.2.1	file_select (compound body)	114
7.2.2	copy_from (compound body)	121
7.2.3	link_from (compound body)	131
7.2.4	perms (compound body)	135
7.2.5	changes (compound body)	137
7.2.6	delete (compound body)	139
7.2.7	rename (compound body)	140
7.2.8	repository	142
7.2.9	edit_line	143
7.2.10	edit_xml	143
7.2.11	edit_defaults (compound body)	143
7.2.12	depth_search (compound body)	144
7.2.13	touch	148

7.2.14	create	149
7.2.15	move_obstructions	149
7.2.16	transformer	150
7.2.17	pathtype	150
7.2.18	acl (compound body)	151
7.3	insert_lines promises in 'edit_line'	152
7.3.1	location (compound body)	153
7.3.2	insert_type	155
7.3.3	insert_select (compound body)	155
7.3.4	expand_scalars	158
7.4	field_edits promises in 'edit_line'	159
7.4.1	edit_field (compound body)	161
7.5	replace_patterns promises in 'edit_line'	164
7.5.1	replace_with (compound body)	165
7.6	delete_lines promises in 'edit_line'	166
7.6.1	not_matching	166
7.6.2	delete_select (compound body)	167
7.7	interfaces promises in 'agent'	170
7.7.1	tcp_ip (compound body)	170
7.8	methods promises in 'agent'	172
7.8.1	usebundle	173
7.9	packages promises in 'agent'	174
7.9.1	install (compound body)	175
7.10	processes promises in 'agent'	175
7.10.1	signals	179
7.10.2	process_stop	180
7.10.3	process_count (compound body)	181
7.10.4	process_select (compound body)	182
7.10.5	restart_class	188
7.11	storage promises in 'agent'	189
7.11.1	mount (compound body)	190
7.11.2	volume (compound body)	193
8	Bundles of server	197
8.1	access promises in 'server'	197
8.1.1	Access Example	198
8.1.2	admit	198
8.1.3	deny	199
8.1.4	maproot	199
8.1.5	ifencrypted	200
8.2	roles promises in 'server'	200
8.2.1	authorize	201

9	Bundles of knowledge	203
9.1	topics promises in ‘knowledge’	203
9.1.1	association (compound body)	204
9.1.2	comment	205
9.2	occurrences promises in ‘knowledge’	206
9.2.1	represents	206
9.2.2	representation	207
10	Special functions	209
10.1	Function accessedbefore	209
10.2	Function accumulated	209
10.3	Function ago	211
10.4	Function canonify	212
10.5	Function changedbefore	212
10.6	Function classmatch	213
10.7	Function execresult	214
10.8	Function fileexists	215
10.9	Function getindices	216
10.10	Function getgid	216
10.11	Function getuid	217
10.12	Function groupexists	218
10.13	Function hash	219
10.14	Function hostrange	220
10.15	Function iprange	220
10.16	Function irange	221
10.17	Function isdir	222
10.18	Function isgreaterthan	222
10.19	Function islessthan	223
10.20	Function islink	224
10.21	Function isnewerthan	225
10.22	Function isplain	226
10.23	Function isvariable	227
10.24	Function now	227
10.25	Function on	228
10.26	Function randomint	228
10.27	Function readfile	229
10.28	Function readintarray	229
10.29	Function readintlist	230
10.30	Function readrealarray	231
10.31	Function readreallist	232
10.32	Function readstringarray	232
10.33	Function readstringlist	233
10.34	Function readtcp	235
10.35	Function regarray	236
10.36	Function regcmp	237
10.37	Function reglist	238
10.38	Function returnszero	238
10.39	Function rrange	239

10.40	Function <code>selectservers</code>	239
10.41	Function <code>strcmp</code>	241
10.42	Function <code>usemodule</code>	241
10.43	Function <code>userexists</code>	242
11	Special Variables	245
11.1	Variable context <code>const</code>	245
11.1.1	Variable <code>const.dollar</code>	245
11.1.2	Variable <code>const.endl</code>	245
11.1.3	Variable <code>const.n</code>	245
11.1.4	Variable <code>const.r</code>	245
11.2	Variable context <code>sys</code>	246
11.2.1	Variable <code>sys.arch</code>	246
11.2.2	Variable <code>sys.cdate</code>	246
11.2.3	Variable <code>sys.class</code>	246
11.2.4	Variable <code>sys.date</code>	247
11.2.5	Variable <code>sys.domain</code>	247
11.2.6	Variable <code>sys.fqhost</code>	247
11.2.7	Variable <code>sys.fstab</code>	247
11.2.8	Variable <code>sys.host</code>	247
11.2.9	Variable <code>sys.interface</code>	248
11.2.10	Variable <code>sys.ipv4</code>	248
11.2.11	Variable <code>sys.ipv4[eth1]</code>	248
11.2.12	Variable <code>sys.ipv4_1[eth1]</code>	248
11.2.13	Variable <code>sys.ipv4_2[eth1]</code>	249
11.2.14	Variable <code>sys.ipv4_3[eth1]</code>	249
11.2.15	Variable <code>sys.long_arch</code>	249
11.2.16	Variable <code>sys.maildir</code>	250
11.2.17	Variable <code>sys.os</code>	250
11.2.18	Variable <code>sys.ostype</code>	250
11.2.19	Variable <code>sys.release</code>	250
11.2.20	Variable <code>sys.resolv</code>	250
11.2.21	Variable <code>sys.uqhost</code>	251
11.2.22	Variable <code>sys.workdir</code>	251

