

Part I. Boost.Build v2 User Manual

Table of Contents

How to use this document	3
Installation	4
Tutorial	6
Hello, world	6
Properties	6
Project Hierarchies	8
Dependent Targets	9
Static and shared libraries	10
Conditions and alternatives	11
Prebuilt targets	11
User documentation	13
Configuration	13
Writing Jamfiles	15
The Build Process	19
Builtin target types	20
Builtin features	24
Differences to Boost.Build V1	25
Extender Manual	27
Introduction	27
Target types	28
Tools and generators	29
Features	32
Main target rules	34
Toolset modules	35
Detailed reference	37
General information	37
Writing Jamfiles	39
Build process	39
Definitions	41
Generators	45
Frequently Asked Questions	47
I'm getting "Duplicate name of actual target" error. What does it mean?	47
Accessing environment variables	47
How to control properties order?	48
How to control the library order on Unix?	48
Can I get output of external program as a variable in a Jamfile?	48
How to get the project-root location?	49
How to change compilation flags for one file?	49
Why are the <code>dll-path</code> and <code>hardcode-dll-paths</code> properties useful?	49
Targets in <code>site-config.jam</code>	50
A. Boost.Build v2 architecture	51
Overview	51
The build layer	51
The tools layer	53
Targets	53

How to use this document

If you've just found out about Boost.Build V2 and want to know if it will work for you, start with *Tutorial*. You can continue with the *User documentation*. When you're ready to try Boost.Build in practice, go to *Installation*.

If you are about to use Boost.Build on your project, or already using it and have a problem, look at *User documentation*.

If you're trying to build a project which uses Boost.Build, look at *Installation* and then read about the section called "Command line".

If you have questions, please post them to our mailing list, and be sure to indicate in the subject line that you're asking about Boost.Build **V2**.

Installation

This section describes how to install Boost.Build from a released Boost source distribution or CVS image.¹ All paths are given relative to the *Boost.Build v2 root directory*, which is located in the `tools/build/v2` subdirectory of a full Boost distribution.²

1. Boost.Build uses Boost.Jam, an extension of the Perforce Jam portable **make** replacement. The recommended way to get Boost.Jam is to **download a prebuilt executable** from SourceForge. If a prebuilt executable is not provided for your platform or you are using Boost's sources in an unreleased state, it may be necessary to build **bjam** from sources included in the Boost source tree.
2. To install Boost.Jam, copy the executable, called **bjam** or **bjam.exe** to a location accessible in your PATH. Go to the Boost.Build root directory and run **bjam --version**. You should see:

```
Boost.Build V2 (Milestone N)
Boost.Jam xx.xx.xx
```

where N is the version of Boost.Build you're using.

3. Configure Boost.Build to recognize the build resources (such as compilers and libraries) you have installed on your system. Open the `user-config.jam` file in the Boost.Build root directory and follow the instructions there to describe your toolsets and libraries, and, if necessary, where they are located.
4. You should now be able to go to the `example/hello/` directory and run **bjam** there. A simple application will be built. You can also play with other projects in the `example/` directory.

If you are using Boost's CVS state, be sure to rebuild **bjam** even if you have a previous version. The CVS version of Boost.Build requires the CVS version of Boost.Jam.

When **bjam** is invoked, it always needs to be able to find the Boost.Build root directory, where the interpreted source code of Boost.Build is located. There are two ways to tell **bjam** about the root directory:

- Set the environment variable `BOOST_BUILD_PATH` to the absolute path of the Boost.Build root directory.
- At the root directory of your project or in any of its parent directories, create a file called `boost-build.jam`, with a single line:

```
boost-build /path/to/boost.build ;
```

¹Note that packages prepared for Unix/Linux systems usually make their own choices about where to put things and even which parts of Boost to include. When we say "released source distribution" we mean a distribution of Boost as released on its SourceForge project page.

²The Boost.Build subset of boost is also distributed separately, for those who are only interested in getting a build tool. The top-level directory of a Boost.Build distribution contains all the subdirectories of the `tools/build/v2` subdirectory from a full Boost distribution, so it is itself a valid Boost.Build root directory. It also contains the `tools/build/jam_src` subdirectory of a full Boost distribution, so you can rebuild Boost.Jam from source.

N.B. When **bjam** is invoked from anywhere in the Boost directory tree *other than* the Boost.Build root and its sub-directories, Boost.Build v1 is used by default. To override the default and use Boost.Build v2, you have to add the `-v2` command line option to all **bjam** invocations.

Tutorial

Hello, world

The simplest project that Boost.Build can construct is stored in `example/hello/` directory. The project is described by a file called `Jamroot` that contains:

```
exe hello : hello.cpp ;
```

Even with this simple setup, you can do some interesting things. First of all, just invoking **bjam** will build the `hello` executable by compiling and linking `hello.cpp`. By default, debug variant is built. Now, to build the release variant of `hello`, invoke

```
bjam release
```

Note that debug and release variants are created in different directories, so you can switch between variants or even build multiple variants at once, without any unnecessary recompilation. Let's extend the example by adding another line to our project's `Jamroot`:

```
exe hello2 : hello.cpp ;
```

Now let us build both the debug and release variants of our project again:

```
bjam debug release
```

Note that two variants of `hello2` are linked. Since we have already built both variants of `hello`, `hello.cpp` won't be recompiled; instead the existing object files will just be linked into the corresponding variants of `hello2`. Now let's remove all the built products:

```
bjam --clean debug release
```

It's also possible to build or clean specific targets. The following two commands, respectively, build or clean only the debug version of `hello2`.

```
bjam hello2  
bjam --clean hello2
```

Properties

To portably represent aspects of target configuration such as debug and release variants, or single- and multi-threaded builds, Boost.Build uses *features* with associated *values*. For example, the `debug-symbols` feature can have a value of `on` or `off`. A *property* is just a (feature, value) pair. When a user initiates a build, Boost.Build automatically translates the requested properties into appropriate command-line flags for invoking toolset components like compilers and linkers.

There are many built-in features that can be combined to produce arbitrary build configurations. The following command builds the project's release variant with inlining disabled and debug symbols enabled:

```
bjam release inlining=off debug-symbols=on
```

Properties on the command-line are specified with the syntax:

```
feature-name=feature-value
```

The `release` and `debug` that we've seen in **bjam** invocations are just a shorthand way to specify values of the `variant` feature. For example, the command above could also have been written this way:

```
bjam variant=release inlining=off debug-symbols=on
```

`variant` is so commonly-used that it has been given special status as an *implicit* feature—Boost.Build will deduce the its identity just from the name of one of its values.

A complete description of features can be found in the section called “Features and properties”.

Build Requests and Target Requirements

The set of properties specified on the command line constitute a *build request*—a description of the desired properties for building the requested targets (or, if no targets were explicitly requested, the project in the current directory). The *actual* properties used for building targets are typically a combination of the build request and properties derived from the project's Jamroot (and its other Jamfiles, as described in the section called “Project Hierarchies”). For example, the locations of `#include` header files are normally not specified on the command-line, but described in Jamfiles as *target requirements* and automatically combined with the build request for those targets. Multithread-enabled compilation is another example of a typical target requirement. The Jamfile fragment below illustrates how these requirements might be specified.

```
exe hello
  : hello.cpp
  : <include>boost <threading>multi
  ;
```

When `hello` is built, the two requirements specified above will always be present. If the build request given on the **bjam** command-line explicitly contradicts a target's requirements, the target requirements usually override (or, in the case of “free” features like `<include>`,¹ augments) the build request.

Tip

The value of the `<include>` feature is relative to the location of Jamroot where it's used.

Project Attributes

If we want the same requirements for our other target, `hello2`, we could simply duplicate them. However, as projects grow, that approach leads to a great deal of repeated boilerplate in Jamfiles. Fortunately, there's a better way. Each project can specify a set of *attributes*, including requirements:

```
project
  : requirements <include>/home/ghost/Work/boost <threading>multi
  ;

exe hello : hello.cpp ;
```

¹ See the section called “Feature Attributes”

```
exe hello2 : hello.cpp ;
```

The effect would be as if we specified the same requirement for both `hello` and `hello2`.

Project Hierarchies

So far we've only considered examples with one project (i.e. with one user-written Boost.Jam file, `Jamroot`). A typical large codebase would be composed of many projects organized into a tree. The top of the tree is called the *project root*. Every subproject is defined by a file called `Jamfile` in a descendant directory of the project root. The parent project of a subproject is defined by the nearest `Jamfile` or `Jamroot` file in an ancestor directory. For example, in the following directory layout:

```
top/
|
+-- Jamroot
|
+-- app/
|   |
|   +-- Jamfile
|   `-- app.cpp
|
|-- util/
|   |
|   +-- foo/
|   .   |
|   .   +-- Jamfile
|   .   `-- bar.cpp
```

the project root is `top/`. Because there is no `Jamfile` in `top/util/`, the projects in `top/app/` and `top/util/foo/` are immediate children of the root project.

Note

When we refer to a “Jamfile,” set in normal type, we mean a file called either `Jamfile` or `Jamroot`. When we need to be more specific, the filename will be set as “`Jamfile`” or “`Jamroot`.”

Projects inherit all attributes (such as requirements) from their parents. Inherited requirements are combined with any requirements specified by the sub-project. For example, if `top/Jamroot` has

```
<include>/home/ghost/local
```

in its requirements, then all of its sub-projects will have it in their requirements, too. Of course, any project can add include paths to those specified by its parents.² More details can be found in the section called “Projects”.

Invoking `bjam` without explicitly specifying any targets on the command-line builds the project rooted in the current directory. Building a project does not automatically cause its sub-projects to be built unless the parent project's `Jamfile` explicitly requests it. In our example, `top/Jamroot` might contain:

```
build-project app ;
```

which would cause the project in `top/app/` to be built whenever the project in `top/` is built. However, targets in `top/util/foo/` will be built only if they are needed by targets in `top/` or `top/app/`.

²Many features will be overridden, rather than added-to, in sub-projects. See the section called “Feature Attributes” for more information

Dependent Targets

Targets that are “needed” by other targets are called *dependencies* of those other targets. The targets that need the other targets are called *dependent* targets.

To get a feeling of target dependencies, let's continue the above example and see how `top/app/Jamfile` can use libraries from `top/util/foo`. If `top/util/foo/Jamfile` contains

```
lib bar : bar.cpp ;
```

then to use this library in `top/app/Jamfile`, we can write:

```
exe app : app.cpp ../util/foo//bar ;
```

While `app.cpp` refers to a regular source file, `../util/foo//bar` is a reference to another target: a library `bar` declared in the Jamfile at `../util/foo`.

Tip

Some other build system have special syntax for listing dependent libraries, for example `LIBS` variable. In Boost.Build, you just add the library to the list of sources.

Suppose we build `app` with:

```
bjam app optimization=full define=USE_ASM
```

Which properties will be used to build `foo`? The answer is that some features are *propagated*—Boost.Build attempts to use dependencies with the same value of propagated features. The `<optimization>` feature is propagated, so both `app` and `foo` will be compiled with full optimization. But `<define>` is not propagated: its value will be added as-is to the compiler flags for `a.cpp`, but won't affect `foo`.

Let's improve this project further. The library probably has some headers that must be used when compiling `app.cpp`. We could manually add the necessary `#include` paths to `app`'s requirements as values of the `<include>` feature, but then this work will be repeated for all programs that use `foo`. A better solution is to modify `util/foo/Jamfile` in this way:

```
project
  : usage-requirements <include>.
  ;
```

```
lib foo : foo.cpp ;
```

Usage requirements are applied not to the target being declared but to its dependents. In this case, `<include>.` will be applied to all targets that directly depend on `foo`.

Another improvement is using symbolic identifiers to refer to the library, as opposed to Jamfile location. In a large project, a library can be used by many targets, and if they all use Jamfile location, a change in directory organization entails much work. The solution is to use project ids—symbolic names not tied to directory layout. First, we need to assign a project id by adding this code to `Jamroot`:

```
use-project /library-example/foo : util/foo ;
```

Second, we modify `app/Jamfile` to use the project id:

```
exe app : app.cpp /library-example/foo//bar ;
```

The `/library-example/foo//bar` syntax is used to refer to the target `bar` in the project with id `/library-example/foo`. We've achieved our goal—if the library is moved to a different directory, only `Jam-root` must be modified. Note that project ids are global—two Jamfiles are not allowed to assign the same project id to different directories.

Tip

If you want all applications in some project to link to a certain library, you can avoid having to specify it directly the sources of every target by using the `<source>` property. For example, if `/boost/filesystem//fs` should be linked to all applications in your project, you can add `<source>/boost/filesystem//fs` to the project's requirements, like this:

```
project
: requirements <source>/boost/filesystem//fs
;
```

Static and shared libraries

Libraries can be either *static*, which means they are included in executable files that use them, or *shared* (a.k.a. *dynamic*), which are only referred to from executables, and must be available at run time. Boost.Build can create and use both kinds.

The kind of library produced from a `lib` target is determined by the value of the `link` feature. Default value is `shared`, and to build static library, the value should be `static`. You can either request static build on the command line:

```
bjam link=static
```

or in the library's requirements:

```
lib l : l.cpp : <link>static ;
```

We can also use the `<link>` property to express linking requirements on a per-target basis. For example, if a particular executable can be correctly built only with the static version of a library, we can qualify the executable's target reference to the library as follows:

```
exe important : main.cpp helpers/<link>static ;
```

No matter what arguments are specified on the **bjam** command-line, `important` will only be linked with the static version of `helpers`.

Specifying properties in target references is especially useful if you use a library defined in some other project (one you can't change) but you still want static (or dynamic) linking to that library in all cases. If that library is used by many targets, you *could* use target references everywhere:

```
exe e1 : e1.cpp /other_project//bar/<link>static ;
exe e10 : e10.cpp /other_project//bar/<link>static ;
```

but that's far from being convenient. A better approach is to introduce a level of indirection. Create a local alias target that refers to the static (or dynamic) version of `foo`:

```
alias foo : /other_project//bar/<link>static ;
exe e1 : e1.cpp foo ;
exe e10 : e10.cpp foo ;
```

The `alias` rule is specifically used to rename a reference to a target and possibly change the properties.

Tip

When one library uses another, you put the second library in the source list of the first. For example:

```
lib utils : utils.cpp /boost/filesystem//fs ;
lib core : core.cpp utils ;
exe app : app.cpp core ;
```

This works no matter what kind of linking is used. When `core` is built as a shared library, it is linked directly into `utils`. Static libraries can't link to other libraries, so when `core` is built as a static library, its dependency on `utils` is passed along to `core`'s dependents, causing `app` to be linked with both `core` and `utils`."

Note

(Note for non-UNIX system). Typically, shared libraries must be installed to a directory in the dynamic linker's search path. Otherwise, applications that use shared libraries can't be started. On Windows, the dynamic linker's search path is given by the `PATH` environment variable. This restriction is lifted when you use `Boost.Build` testing facilities—the `PATH` variable will be automatically adjusted before running executable.

Conditions and alternatives

Sometimes, particular relationships need to be maintained among a target's build properties. This can be achieved with *conditional requirement*. For example, you might want to set specific `#defines` when a library is built as shared, or when a target's `release` variant is built in release mode.

```
lib network : network.cpp
  : <link>shared:<define>NETWORK_LIB_SHARED
  <variant>release:<define>EXTRA_FAST
  ;
```

In the example above, whenever `network` is built with `<link>shared`, `<define>NETWORK_LIB_SHARED` will be in its properties, too.

Sometimes the ways a target is built are so different that describing them using conditional requirements would be hard. For example, imagine that a library actually uses different source files depending on the toolset used to build it. We can express this situation using *target alternatives*:

```
lib demangler : dummy_demangler.cpp ; # alternative 1
lib demangler : demangler_gcc.cpp : <toolset>gcc ; # alternative 2
lib demangler : demangler_msvc.cpp : <toolset>msvc ; # alternative 3
```

In the example above, when built with `gcc` or `msvc`, `demangler` will use a source file specific to the toolset. Otherwise, it will use a generic source file, `dummy_demangler.cpp`.

Prebuilt targets

To link to libraries whose build instructions aren't given in a Jamfile, you need to create `lib` targets with an appro-

appropriate file property. Target alternatives can be used to associate multiple library files with a single conceptual target. For example:

```
# util/lib2/Jamfile
lib lib2
:
: <file>lib2_release.a <variant>release
;

lib lib2
:
: <file>lib2_debug.a <variant>debug
;
```

This example defines two alternatives for `lib2`, and for each one names a prebuilt file. Naturally, there are no sources. Instead, the `<file>` feature is used to specify the file name.

Once a prebuilt target has been declared, it can be used just like any other target:

```
exe app : app.cpp ../util/lib2//lib2 ;
```

As with any target, the alternative selected depends on the properties propagated from `lib2`'s dependents. If we build the the release and debug versions of `app` will be linked with `lib2_release.a` and `lib2_debug.a`, respectively.

System libraries—those that are automatically found by the toolset by searching through some set of predetermined paths—should be declared almost like regular ones:

```
lib pythonlib : : <name>python22 ;
```

We again don't specify any sources, but give a name that should be passed to the compiler. If the `gcc` toolset were used to link an executable target to `pythonlib`, `-lpython22` would appear in the command line (other compilers may use different options).

We can also specify where the toolset should look for the library:

```
lib pythonlib : : <name>python22 <search>/opt/lib ;
```

And, of course, target alternatives can be used in the usual way:

```
lib pythonlib : : <name>python22 <variant>release ;
lib pythonlib : : <name>python22_d <variant>debug ;
```

A more advanced use of prebuilt targets is described in the section called “Targets in `site-config.jam`”.

User documentation

This section will provide the information necessary to create your own projects using Boost.Build. The information provided here is relatively high-level, and *Detailed reference* as well as the on-line help system must be used to obtain low-level documentation (see ???).

Boost.Build actually consists of two parts - Boost.Jam, a build engine with its own interpreted language, and Boost.Build itself, implemented in Boost.Jam's language. The chain of events when you type **bjam** on the command line is:

1. Boost.Jam tries to find Boost.Build and loads the top-level module. The exact process is described in the section called "Initialization"
2. The top-level module loads user-defined configuration files, `user-config.jam` and `site-config.jam`, which define available toolsets.
3. The Jamfile in the current directory is read. That in turn might cause reading of further Jamfiles. As a result, a tree of projects is created, with targets inside projects.
4. Finally, using the build request specified on the command line, Boost.Build decides which targets should be built, and how. That information is passed back to Boost.Jam, which takes care of actually running commands.

So, to be able to successfully use Boost.Build, you need to know only three things:

- How to configure Boost.Build
- How to write Jamfiles
- How the build process works
- Some Basics about the Boost.Jam language. See the Boost.Jam and Classic Jam documentation.

Configuration

The Boost.Build configuration is specified in the file `user-config.jam`. You can edit the one that comes with Boost.Build, or create a copy in your home directory and edit that. (See the reference for the exact search paths.) The primary function of that file is to declare which compilers and other tools are available. The simplest syntax to configure a tool is:

```
using tool-name ;
```

The `using` rule is given a name of tool, and will make that tool available to Boost.Build. For example, `using gcc ;` will make the `gcc` compiler available.

Since nothing but a tool name is specified, Boost.Build will pick some default settings. For example, it will use the `gcc` executable found in the `PATH`, or look in some known installation locations. In most cases, this strategy works automatically. In case you have several versions of a compiler, it's installed in some unusual location, or you need to tweak its configuration, you'll need to pass additional parameters to the `using` rule. The parameters to `using` can be different for each tool. You can obtain specific documentation for any tool's configuration parameters by invoking

```
bjam --help tool-name.init
```

That said, for all the compiler toolsets Boost.Build supports out-of-the-box, the list of parameters to `using` is the same: *toolset-name*, *version*, *invocation-command*, and *options*.

The *version* parameter identifies the toolset version, in case you have several installed. It can have any form you like, but it's recommended that you use a numeric identifier like 7.1.

The *invocation-command* parameter is the command that must be executed to run the compiler. This parameter can usually be omitted if the compiler executable

- has its “usual name” and is in the `PATH`, or
- was installed in a standard “installation directory”, or
- can be found through a global mechanism like the Windows registry.

For example:

```
using msvc : 7.1 ;
using gcc ;
```

If the compiler can be found in the `PATH` but only by a nonstandard name, you can just supply that name:

```
using gcc : : g++-3.2 ;
```

Otherwise, it might be necessary to supply the complete path to the compiler executable:

```
using msvc : : Z:/Programs/Microsoft Visual Studio/vc98/bin/cl ;
```

Some Boost.Build toolsets will use that path to take additional actions required before invoking the compiler, such as calling vendor-supplied scripts to set up its required environment variables.

To configure several versions of a toolset, simply invoke the `using` rule multiple times:

```
using gcc : 3.3 ;
using gcc : 3.4 : g++-3.4 ;
using gcc : 3.2 : g++-3.2 ;
```

Note that in the first call to `using`, the compiler found in the `PATH` will be used, and there's no need to explicitly specify the command.

As shown above, both the *version* and *invocation-command* parameters are optional, but there's an important restriction: if you configure the same toolset more than once, you must pass the *version* parameter every time. For example, the following is not allowed:

```
using gcc ;
using gcc : 3.4 : g++-3.4 ;
```

because the first `using` call does not specify a *version*.

The *options* parameter is used to fine-tune the configuration. All of Boost.Build's standard compiler toolsets accept properties of the four builtin features `cflags`, `cxxflags`, `compileflags` and `linkflags` as *options* specifying flags that will be always passed to the corresponding tools. Values of the `cflags` feature are passed dir-

ectly to the C compiler, values of the `cxxflags` feature are passed directly to the C++ compiler, and values of the `compileflags` feature are passed to both. For example, to configure a `gcc` toolset so that it always generates 64-bit code you could write:

```
using gcc : 3.4 : : <compileflags>-m64 <linkflags>-m64 ;
```

Writing Jamfiles

Overview

Jamfiles are the thing that is most important to the user, because they declare the targets that should be built. Jamfiles are also used for organizing targets—each Jamfile is a separate project that can be built independently from the other projects.

Jamfiles mostly contain calls to Boost.Build functions that do all the work, specifically:

- declare main targets
- define project properties
- do various other things

Main targets

A *Main target* is a user-defined named entity that can be built, for example an executable file. Declaring a main target is usually done using one of the main target rules described in the section called “Builtin target types”. The user can also declare custom main target rules as shown in the section called “Main target rules”.

Most main target rules in Boost.Build can be invoked with a common syntax:

```
rule-name main-target-name
  : sources...
  : requirements...
  : default-build...
  : usage-requirements...
  ;
```

- *main-target-name* is the name used to request the target on command line and to use it from other main targets. A main target name may contain alphanumeric characters, dashes (‘-’), and underscores (‘_’).
- *sources* is the list of source files and other main targets that must be combined.
- *requirements* is the list of properties that must always be present when this main target is built.
- *default-build* is the list of properties that will be used unless some other value of the same feature is already specified, e.g. on the command line or by propagation from a dependent target.
- *usage-requirements* is the list of properties that will be propagated to all main targets that use this one, i.e. to all its dependents.

Some main target rules have a shorter list of parameters; consult their documentation for details.

Note that the actual requirements, default-build and usage-requirements attributes for a target are obtained by combining the explicitly specified ones with those specified for the project where a target is declared.

The list of sources specifies what should be processed to get the resulting targets. Most of the time, it's just a list of files. Sometimes, you'll want to automatically construct the list of source files rather than having to spell it out manually, in which case you can use the `glob` rule. Here are two examples:

```
exe a : a.cpp ;           # a.cpp is the only source file
exe b : [ glob *.cpp ] ; # all .cpp files in this directory are sources
```

Unless you specify a files with absolute path, the name is considered relative to the source directory -- which is typically the directory where the Jamfile is located, but can be changed as described in the section called “Projects” [?].

The list of sources can also refer to other main targets. Targets in the same project can be referred to by name, while targets in other projects need to be qualified with a directory or a symbolic project name. For example:

```
lib helper : helper.cpp ;
exe a : a.cpp helper ;
exe b : b.cpp ../utils ;
exe c : c.cpp /boost/program_options//program_options ;
```

The first exe uses the library defined in the same project. The second one uses some target (most likely library) defined by Jamfile one level higher. Finally, the third target uses some C++ Boost library, referring to it by its absolute symbolic name. More information about it can be found in the section called “Dependent Targets” and the section called “Target identifiers and references”.

Requirements are the properties that should always be present when building a target. Typically, they are includes and defines:

```
exe hello : hello.cpp : <include>/opt/boost <define>MY_DEBUG ;
```

In special circumstances, other properties can be used, for example if a library can only be built statically, or a file can't be compiled with optimization due to a compiler bug, one can use

```
lib util : util.cpp : <link>static ;
obj main : main.cpp : <optimization>off ;
```

Sometimes requirements are necessary only for a specific compiler or build variant. Conditional properties can be used in that case:

```
lib util : util.cpp : <toolset>msvc:<link>static ;
```

Whenever `<toolset>msvc` property is in build properties, the `<link>static` property will be included as well. Conditional requirements can be “chained”:

```
lib util : util.cpp : <toolset>msvc:<link>static
                    <link>static:<define>STATIC_LINK ;
```

will set of static link and the `STATIC_LINK` define on the `msvc` toolset.

The `default-build` parameter is a set of properties to be used if the build request does not otherwise specify a value for features in the set. For example:


```
exe hello : hello.cpp : : <threading>multi ;
```

would build a multi-threaded target in unless the user explicitly requests a single-threaded version. The difference between requirements and default-build is that requirements cannot be overridden in any way.

A target of the same name can be declared several times, in which case each declaration is called an *alternative*. When the target is built, one of the alternatives will be selected and used. Alternatives need not be defined by the same main target rule. For example,

```
lib helpers : helpers.hpp ; # a header-only library
alias helpers : helpers.lib : <toolset>msvc ; # except on msvc
```

The actual commands used to build any given main target can differ greatly from platform to platform. For example, you might have different lists of sources for different compilers, or different options for those compilers. Two approaches to this are explained in the tutorial.

Sometimes a main target is really needed only by some other main target. For example, a rule that declares a test-suite uses a main target that represent test, but those main targets are rarely needed by themselves.

It is possible to declare a target inline, i.e. the "sources" parameter may include calls to other main rules. For example:

```
exe hello : hello.cpp
  [ obj helpers : helpers.cpp : <optimization>off ] ;
```

Will cause "helpers.cpp" to be always compiled without optimization. When referring to an inline main target, its declared name must be prefixed by its parent target's name and two dots. In the example above, to build only helpers, one should run `bjam hello..helpers`.

Projects

As mentioned before, targets are grouped into projects, and each Jamfile is a separate project. Projects are useful because they allow us to group related targets together, define properties common to all those targets, and assign a symbolic name to the project that can be used in referring to its targets.

Projects are named using the `project` rule, which has the following syntax:

```
project id : attributes ;
```

Here, *attributes* is a sequence of rule arguments, each of which begins with an attribute-name and is followed by any number of build properties. The list of attribute names along with its handling is also shown in the table below. For example, it is possible to write:

```
project tennis
  : requirements <threading>multi
  : default-build release
  ;
```

The possible attributes are listed below.

Project id is a short way to denote a project, as opposed to the Jamfile's pathname. It is a hierarchical path, unrelated to filesystem, such as "boost/thread". Target references make use of project ids to specify a target.

Source location specifies the directory where sources for the project are located.

Project requirements are requirements that apply to all the targets in the projects as well as all subprojects.

Default build is the build request that should be used when no build request is specified explicitly.

The default values for those attributes are given in the table below.

Table 1.

Attribute	Name for the 'project' rule	Default value	Handling by the 'project' rule
Project id	none	none	Assigned from the first parameter of the 'project' rule. It is assumed to denote absolute project id.
Source location	source-location	The location of jamfile for the project	Sets to the passed value
Requirements	requirements	The parent's requirements	The parent's requirements are refined with the passed requirement and the result is used as the project requirements.
Default build	default-build	none	Sets to the passed value
Build directory	build-dir	Empty if the parent has no build directory set. Otherwise, the parent's build directory with with the relative path from parent to the current project appended to it.	Sets to the passed value, interpreted as relative to the project's location.

Besides defining projects and main targets, Jamfiles commonly invoke utility rules such as `constant` and `path-constant`, which inject a specified Boost.Jam variable setting into this project's Jamfile module and those of all its subprojects. See the section called “Jamfile Utility Rules” for a complete description of these utility rules. Jamfiles are regular Boost.Jam source files and Boost.Build modules, so naturally they can contain any kind of Boost.Jam code, including rule definitions.

Each subproject inherits attributes, constants and rules from its parent project, which is defined by the nearest Jamfile in an ancestor directory above the subproject. The top-level project is declared in a file called `Jamroot` rather than `Jamfile`. When loading a project, Boost.Build looks for either `Jamroot` or `Jamfile`. They are handled identically, except that if the file is called `Jamroot`, the search for a parent project is not performed.

Even when building in a subproject directory, parent project files are always loaded before those of their subprojects, so that every definition made in a parent project is always available to its children. The loading order of any other projects is unspecified. Even if one project refers to another via `use-project`, or a target reference, no specific order should be assumed.

Note

Giving the root project the special name “`Jamroot`” ensures that Boost.Build won't misinterpret a directory above it as the project root just because the directory contains a Jamfile.

Jamfile Utility Rules

The following table describes utility rules that can be used in Jamfiles. Detailed information for any of these rules can be obtained by running:

```
bjam --help project.rulename
```

Table 2.

Rule	Semantics
project	Define this project's symbolic ID or attributes.
use-project	Make another project known so that it can be referred to by symbolic ID.
build-project	Cause another project to be built when this one is built.
explicit	State that a target should be built only by explicit request.
glob	Translate a list of shell-style wildcards into a corresponding list of files.
constant	Injects a variable setting into this project's Jamfile module and those of all its subprojects.
path-constant	Injects a variable set to a path value into this project's Jamfile module and those of all its subprojects. If the value is a relative path it will be adjusted for each subproject so that it refers to the same directory.

The Build Process

When you've described your targets, you want Boost.Build to run the right tools and create the needed targets. This section will describe two things: how you specify what to build, and how the main targets are actually constructed.

The most important thing to note is that in Boost.Build, unlike other build tools, the targets you declare do not correspond to specific files. What you declare in a Jamfile is more like a “metatarget.” Depending on the properties you specify on the command line, each metatarget will produce a set of real targets corresponding to the requested properties. It is quite possible that the same metatarget is built several times with different properties, producing different files.

Tip

This means that for Boost.Build, you cannot directly obtain a build variant from a Jamfile. There could be several variants requested by the user, and each target can be built with different properties.

Build request

The command line specifies which targets to build and with which properties. For example:

```
bjam app1 lib1//lib1 toolset=gcc variant=debug optimization=full
```

would build two targets, "app1" and "lib1/lib1" with the specified properties. You can refer to any targets, using target id and specify arbitrary properties. Some of the properties are very common, and for them the name of the property can be omitted. For example, the above can be written as:

```
bjam app1 lib1//lib1 gcc debug optimization=full
```

The complete syntax, which has some additional shortcuts, is described in the section called “Command line”.

Building a main target

When you request, directly or indirectly, a build of a main target with specific requirements, the following steps are made. Some brief explanation is provided, and more details are given in the section called “Build process”.

1. Applying default build. If the default-build property of a target specifies a value of a feature that is not present in the build request, that value is added.
2. Selecting the main target alternative to use. For each alternative we look how many properties are present both in alternative's requirements, and in build request. The alternative with large number of matching properties is selected.
3. Determining "common" properties. The build request is refined with target's requirements. The conditional properties in requirements are handled as well. Finally, default values of features are added.
4. Building targets referred by the sources list and dependency properties. The list of sources and the properties can refer to other target using target references. For each reference, we take all propagated properties, refine them by explicit properties specified in the target reference, and pass the resulting properties as build request to the other target.
5. Adding the usage requirements produced when building dependencies to the "common" properties. When dependencies are built in the previous step, they return both the set of created "real" targets, and usage requirements. The usage requirements are added to the common properties and the resulting property set will be used for building the current target.
6. Building the target using generators. To convert the sources to the desired type, Boost.Build uses "generators" - -- objects that correspond to tools like compilers and linkers. Each generator declares what type of targets it can produce and what type of sources it requires. Using this information, Boost.Build determines which generators must be run to produce a specific target from specific sources. When generators are run, they return the "real" targets.
7. Computing the usage requirements to be returned. The conditional properties in usage requirements are expanded and the result is returned.

Building a project

Often, a user builds a complete project, not just one main target. In fact, invoking **bjam** without arguments builds the project defined in the current directory.

When a project is built, the build request is passed without modification to all main targets in that project. It's possible to prevent implicit building of a target in a project with the `explicit` rule:

```
explicit hello_test ;
```

would cause the `hello_test` target to be built only if explicitly requested by the user or by some other target.

The Jamfile for a project can include a number of `build-project` rule calls that specify additional projects to be built.

Builtin target types

Programs

Programs are created using the `exe` rule, which follows the common syntax. For example:

```
exe hello : hello.cpp some_library.lib /some_project//library
          : <threading>multi
          ;
```

This will create an executable file from the sources -- in this case, one C++ file, one library file present in the same directory, and another library that is created by Boost.Build. Generally, sources can include C and C++ files, object files and libraries. Boost.Build will automatically try to convert targets of other types.

Tip

On Windows, if an application uses dynamic libraries, and both the application and the libraries are built by Boost.Build, it's not possible to immediately run the application, because the `PATH` environment variable should include the path to the libraries. It means you have to either add the paths manually, or place the application and the libraries to the same directory, for example using the stage rule.

Libraries

Libraries are created using the `lib` rule, which follows the common syntax. For example:

```
lib helpers : helpers.cpp : <include>boost : : <include>. ;
```

In the most common case, the `lib` creates a library from the specified sources. Depending on the value of `<link>` feature the library will be either static or shared. There are two other cases. First is when the library is installed somewhere in compiler's search paths, and should be searched by the compiler (typically, using the `-l` option). The second case is where the library is available as a prebuilt file and the full path is known.

The syntax for these case is given below:

```
lib z : : <name>z <search>/home/ghost ;
lib compress : : <file>/opt/libs/compress.a ;
```

The `name` property specifies the name that should be passed to the `-l` option, and the `file` property specifies the file location. The `search` feature specifies paths in which to search for the library. That feature can be specified several times, or it can be omitted, in which case only default compiler paths will be searched.

The difference between using the `file` feature as opposed to the `name` feature together with the `search` feature is that `file` is more precise. A specific file will be used. On the other hand, the `search` feature only adds a library path, and the `name` feature gives the basic name of the library. The search rules are specific to the linker. For example, given these definition:

```
lib a : : <variant>release <file>/pool/release/a.so ;
lib a : : <variant>debug <file>/pool/debug/a.so ;
lib b : : <variant>release <file>/pool/release/b.so ;
lib b : : <variant>debug <file>/pool/debug/b.so ;
```

It's possible to use release version of `a` and debug version of `b`. Had we used the `name` and `search` features, the linker would always pick either release or debug versions.

For convenience, the following syntax is allowed:

```
lib z ;
```

```
lib gui db aux ;
```

and is does exactly the same as:

```
lib z : : <name>z ;  
lib gui : : <name>gui ;  
lib db : : <name>db ;  
lib aux : : <name>aux ;
```

When a library uses another library you should put that another library in the list of sources. This will do the right thing in all cases. For portability, you should specify library dependencies even for searched and prebuilt libraries, otherwise, static linking on Unix won't work. For example:

```
lib z ;  
lib png : z : <name>png ;
```

Note

When a library (say, *a*), that has another library, (say, *b*) is linked dynamically, the *b* library will be incorporated in *a*. (If *b* is dynamic library as well, then *a* will only refer to it, and not include any extra code.) When the *a* library is linked statically, Boost.Build will assure that all executables that link to *a* will also link to *b*.

One feature of Boost.Build that is very important for libraries is usage requirements. For example, if you write:

```
lib helpers : helpers.cpp : : : <include>. ;
```

then the compiler include path for all targets that use *helpers* will contain the directory where the target is defined.path to "helpers.cpp". The user only needs to add *helpers* to the list of sources, and needn't consider the requirements its use imposes on a dependent target. This feature greatly simplifies Jamfiles.

Note

If you don't want shared libraries to include all libraries that are specified in sources (especially statically linked ones), you'd need to use the following:

```
lib b : a.cpp ;  
lib a : a.cpp : <use>b : : <library>b ;
```

This specifies that *a* uses *b*, and causes all executables that link to *a* also link to *b*. In this case, even for shared linking, the *a* library won't even refer to *b*.

Alias

The *alias* rule follows the common syntax. For example:

```
alias core : im reader writer ;
```

will build the sources and return the generated source targets without modification.

The *alias* rule is a convenience tool. If you often build the same group of targets at the same time, you can define an alias to save typing.

Another use of the *alias* rule is to change build properties. For example, if you always want static linking for a specific C++ Boost library, you can write the following:

```
alias threads : /boost/thread//boost_thread : <link>static ;
```

and use only the `threads` alias in your Jamfiles.

You can also specify usage requirements for the `alias` target. If you write the following:

```
alias header_only_library : : : : <include>/usr/include/header_only_library ;
```

then using `header_only_library` in sources will only add an include path. Also note that when there are some sources, their usage requirements are propagated, too. For example:

```
lib lib : lib.cpp : : : <include>. ;
alias lib_alias ;
exe main : main.cpp lib_alias ;
```

will compile `main.cpp` with the additional include.

Installing

For installing a built target you should use the `install` rule, which follows the common syntax. For example:

```
install dist : hello helpers ;
```

will cause the targets `hello` and `helpers` to be moved to the `dist` directory, relative to Jamfile's directory. The directory can be changed with the `location` property:

```
install dist : hello helpers : <location>/usr/bin ;
```

While you can achieve the same effect by changing the target name to `/usr/bin`, using the `location` property is better, because it allows you to use a mnemonic target name.

The `location` property is especially handy when the location is not fixed, but depends on build variant or environment variables:

```
install dist : hello helpers : <variant>release:<location>dist/release
                             <variant>debug:<location>dist/debug ;
install dist2 : hello helpers : <location>$(DIST) ;
```

See also conditional properties and environment variables

Specifying the names of all libraries to install can be boring. The `install` allows you to specify only the top-level executable targets to install, and automatically install all dependencies:

```
install dist : hello
              : <install-dependencies>on <install-type>EXE
              : <install-type>LIB
              ;
```

will find all targets that `hello` depends on, and install all of the which are either executables or libraries. More specifically, for each target, other targets that were specified as sources or as dependency properties, will be recursively found. One exception is that targets referred with the `use` feature are not considered, because that feature is typically used to refer to header-only libraries. If the set of target types is specified, only targets of that type will be installed, otherwise, all found target will be installed.

The `alias` rule can be used when targets must be installed into several directories:

```
install install : install-bin install-lib ;
install install-bin : applications : /usr/bin ;
install install-lib : helper : /usr/lib ;
```

Because the `install` rule just copies targets, most free features¹ have no effect when used in requirements of the `install`. The only two which matter are `dependency` and, on Unix, `dll-path`.

Note

(Unix specific). On Unix, executables built with Boost.Build typically contain the list of paths to all used dynamic libraries. For installing, this is not desired, so Boost.Build relinks the executable with an empty list of paths. You can also specify additional paths for installed executables with the `dll-path` feature.

Testing

Boost.Build has convenient support for running unit tests. The simplest way is the `unit-test` rule, which follows the common syntax. For example:

```
unit-test helpers_test : helpers_test.cpp helpers ;
```

The `unit-test` rule behaves like the `exe` rule, but after the executable is created it is run. If the executable returns an error code, the build system will also return an error and will try running the executable on the next invocation until it runs successfully. This behaviour ensures that you can't miss a unit test failure.

There are rules for more elaborate testing: `compile`, `compile-fail`, `run` and `run-fail`. They are more suitable for automated testing, and are not covered here.

Builtin features

`variant`

A feature that combines several low-level features, making it easy to request common build configurations.

Allowed values: `debug`, `release`, `profile`

The value `debug` expands to

```
<optimization>off <debug-symbols>on <inlining>off <runtime>
```

The value `release` expands to

```
<optimization>speed <debug-symbols>off <inlining>full <runtime>
```

The value `profile` expands to the same as `release`, plus:

¹see the definition of "free" in the section called "Feature Attributes".

`<profiling>on <debug-symbols>on`

Rationale: Runtime debugging is on in debug builds to suit the expectations of people used to various IDEs. It's assumed other folks don't have any specific expectation in this point.

link	A feature that controls how libraries are built. Allowed values: <code>shared</code> , <code>static</code>
source	The <code><source>X</code> feature has the same effect on building a target as putting <code>X</code> in the list of sources. The feature is sometimes more convenient: you can put <code><source>X</code> in the requirements for a project and <code>X</code> will be included as a source in all of the project's main targets.
library	This feature is equivalent to the <code><source></code> feature, and exists for backward compatibility reasons.
dependency	Introduces a dependency on the target named by the value of this feature (so it will be brought up-to-date whenever the target being declared is), and adds its usage requirements to the build properties of the target being declared. The dependency is not used in any other way. The primary use case is when you want the usage requirements (such as <code>#include</code> paths) of some library to be applied, but don't want to link to it.
use	Introduces a dependency on the target named by the value of this feature (so it will be brought up-to-date whenever the target being declared is), and adds its usage requirements to the build properties of the target being declared. The dependency is not used in any other way. The primary use case is when you want the usage requirements (such as <code>#include</code> paths) of some library to be applied, but don't want to link to it.
dll-path	Specify an additional directory where the system should look for shared libraries when the executable or shared library is run. This feature only affects Unix compilers. Please see the section called “Why are the <code>dll-path</code> and <code>hardcode-dll-paths</code> properties useful?” in <i>Frequently Asked Questions</i> for details.
hardcode-dll-paths	Controls automatic generation of <code>dll-path</code> properties. Allowed values: <code>true</code> , <code>false</code> . This property is specific to Unix systems. If an executable is built with <code><hardcode-dll-paths>true</code> , the generated binary will contain the list of all the paths to the used shared libraries. As the result, the executable can be run without changing system paths to shared libraries or installing the libraries to system paths. This is very convenient during development. Please see the FAQ entry for details. Note that on Mac OSX, the paths are unconditionally hardcoded by the linker, and it's not possible to disable that behaviour.
<code>cflags</code> , <code>cxxflags</code> , <code>linkflags</code>	The value of those features is passed without modification to the corresponding tools. For <code>cflags</code> that's both C and C++ compilers, for <code>cxxflags</code> that's C++ compiler and for <code>linkflags</code> that's linker. The features are handy when you're trying to do something special that cannot be achieved by higher-level feature in Boost.Build.

Differences to Boost.Build V1

While Boost.Build V2 is based on the same ideas as Boost.Build V1, some of the syntax was changed, and some

new important features were added. This chapter describes most of the changes.

Configuration

In V1, there were two methods to configure a toolset. One was to set some environment variable, or use the `-s` command line option to set a variable inside BJam. Another method was to create a new toolset module that would set the variables and then invoke the base toolset. Neither method is necessary now: the `using` rule provides a consistent way to initialize a toolset, including several versions. See the section called “Configuration” for details.

Writing Jamfiles

Probably one of the most important differences in V2 Jamfiles is the use of project requirements. In V1, if several targets had the same requirements (for example, a common `#include` path), it was necessary to manually write the requirements or use a helper rule or template target. In V2, the common properties can be specified with the `requirements` project attribute, as documented in the section called “Projects”.

Usage requirements also help to simplify Jamfiles. If a library requires all clients to use specific `#include` paths or macros when compiling code that depends on the library, that information can be cleanly represented.

The difference between `lib` and `dll` targets in V1 is completely eliminated in V2. There's only one library target type, `lib`, which can create either static or shared libraries depending on the value of the `<link>` feature. If your target should be only built in one way, you can add `<link>shared` or `<link>static` to its requirements.

The syntax for referring to other targets was changed a bit. While in V1 one would use:

```
exe a : a.cpp <lib>../foo/bar ;
```

the V2 syntax is:

```
exe a : a.cpp ../foo//bar ;
```

Note that you don't need to specify the type of other target, but the last element should be separated from the others by a double slash to indicate that you're referring to target `bar` in project `../foo`, and not to project `../foo/bar`.

Build process

The command line syntax in V2 is completely different. For example

```
bjam -sTOOLS=msvc -sBUILD=release some_target
```

now becomes:

```
bjam toolset=msvc variant=release some_target
```

or, using implicit features, just:

```
bjam msvc release some_target
```

See the reference for a complete description of the syntax.

Extender Manual

Introduction

This document explains how to extend Boost.Build to accomodate your local requirements. Let's start with a simple but realistic example.

Say you're writing an application that generates C++ code. If you ever did this, you know that it's not nice. Embedding large portions of C++ code in string literals is very awkward. A much better solution is:

1. Write the template of the code to be generated, leaving placeholders at the points that will change
2. Access the template in your application and replace placeholders with appropriate text.
3. Write the result.

It's quite easy to achieve. You write special verbatim files that are just C++, except that the very first line of the file contains the name of a variable that should be generated. A simple tool is created that takes a verbatim file and creates a cpp file with a single `char*` variable whose name is taken from the first line of the verbatim file and whose value is the file's properly quoted content.

Let's see what Boost.Build can do.

First off, Boost.Build has no idea about "verbatim files". So, you must register a new target type. The following code does it:

```
import type ;
type.register VERBATIM : vrb ;
```

The first parameter to `type.register` gives the name of the declared type. By convention, it's uppercase. The second parameter is the suffix for files of this type. So, if Boost.Build sees `code.vrb` in a list of sources, it knows that it's of type VERBATIM.

Next, you tell Boost.Build that the verbatim files can be transformed into C++ files in one build step. A *generator* is a template for a build step that transforms targets of one type (or set of types) into another. Our generator will be called `verbatim.inline-file`; it transforms VERBATIM files into CPP files:

```
import generators ;
generators.register-standard verbatim.inline-file : VERBATIM : CPP ;
```

Lastly, you have to inform Boost.Build about the shell commands used to make that transformation. That's done with an `actions` declaration.

```
actions inline-file
{
    "./inline-file.py" $(<) $(>)
}
```

Now, we're ready to tie it all together. Put all the code above in file `verbatim.jam`, add `import verbatim ;` to `project-root.jam`, and it's possible to write the following in Jamfile:

```
exe codegen : codegen.cpp class_template.verbatim usage.verbatim ;
```

The verbatim files will be automatically converted into C++ and linked it.

In the subsequent sections, we will extend this example, and review all the mechanisms in detail. The complete code is available in `example/customization` directory.

Target types

The first thing we did in the introduction was declaring a new target type:

```
import type ;
type.register VERBATIM : verbatim ;
```

The `type` is the most important property of a target. Boost.Build can automatically generate necessary build actions only because you specify the desired type (using the different main target rules), and because Boost.Build can guess the type of sources from their extensions.

The first two parameters for the `type.register` rule are the name of new type and the list of extensions associated with it. A file with an extension from the list will have the given target type. In the case where a target of the declared type is generated from other sources, the first specified extension will be used.

Sometimes you want to change the suffix used for generated targets depending on build properties, such as `toolset`. For example, some compiler uses extension `elf` for executable files. You can use the `type.set-generated-target-suffix` rule:

```
type.set-generated-target-suffix EXE : <toolset>elf : elf ;
```

A new target type can be inherited from an existing one.

```
type.register PLUGIN : : SHARED_LIB ;
```

The above code defines a new type derived from `SHARED_LIB`. Initially, the new type inherits all the properties of the base type - in particular generators and suffix. Typically, you'll change the new type in some way. For example, using `type.set-generated-target-suffix` you can set the suffix for the new type. Or you can write special generator for the new type. For example, it can generate additional meta-information for plugin. In either way, the `PLUGIN` type can be used whenever `SHARED_LIB` can. For example, you can directly link plugins to an application.

A type can be defined as "main", in which case Boost.Build will automatically declare a main target rule for building targets of that type. More details can be found later.

Scanners

Sometimes, a file can refer to other files via some include mechanism. To make Boost.Build track dependencies to the included files, you need to provide a scanner. The primary limitation is that only one scanner can be assigned to a target type.

First, we need to declare a new class for the scanner:

```
class verbatim-scanner : common-scanner
```

```
{
  rule pattern ( )
  {
    return "//###include[ ]*\\"([^\"]*)\\" ;
  }
}
```

All the complex logic is in the `common-scanner` class, and you only need to override the method that returns the regular expression to be used for scanning. The paranthesis in the regular expression indicate which part of the string is the name of the included file. Only the first parenthesized group in the regular expression will be recognized; if you can't express everything you want that way, you can return multiple regular expressions, each of which contains a parenthesized group to be matched.

After that, we need to register our scanner class:

```
scanner.register verbatim-scanner : include ;
```

The value of the second parameter, in this case `include`, specifies the properties that contain the list of paths that should be searched for the included files.

Finally, we assign the new scanner to the `VERBATIM` target type:

```
type.set-scanner VERBATIM : verbatim-scanner ;
```

That's enough for scanning include dependencies.

Tools and generators

This section will describe how Boost.Build can be extended to support new tools.

For each additional tool, a Boost.Build object called generator must be created. That object has specific types of targets that it accepts and produces. Using that information, Boost.Build is able to automatically invoke the generator. For example, if you declare a generator that takes a target of the type `D` and produces a target of the type `OBJ`, when placing a file with extension `.d` in a list of sources will cause Boost.Build to invoke your generator, and then to link the resulting object file into an application. (Of course, this requires that you specify that the `.d` extension corresponds to the `D` type.)

Each generator should be an instance of a class derived from the `generator` class. In the simplest case, you don't need to create a derived class, but simply create an instance of the `generator` class. Let's review the example we've seen in the introduction.

```
import generators ;
generators.register-standard verbatim.inline-file : VERBATIM : CPP ;
actions inline-file
{
  "./inline-file.py" $(<) $(>)
}
```

We declare a standard generator, specifying its id, the source type and the target type. When invoked, the generator will create a target of type `CPP` with a source target of type `VERBATIM` as the only source. But what command will be used to actually generate the file? In `bjam`, actions are specified using named "actions" blocks and the name of the action block should be specified when creating targets. By convention, generators use the same name of the action block as their own id. So, in above example, the "inline-file" actions block will be used to convert the source into the target.

There are two primary kinds of generators: standard and composing, which are registered with the `generators.register-standard` and the `generators.register-composing` rules, respectively. For example:

```
generators.register-standard verbatim.inline-file : VERBATIM : CPP ;
generators.register-composing mex.mex : CPP LIB : MEX ;
```

Standard generators take a *single* source of type VERBATIM and produces a result. The second generator takes any number of sources, which can have either the CPP or the LIB type. Composing generators are typically used for generating top-level target type. For example, the first generator invoked when building an exe target is a composing generator corresponding to the proper linker.

You should also know about two specific function for registering generators: `generators.register-c-compiler` and `generators.register-linker`. The first sets up header dependency scanning for C files, and the seconds handles various complexities like searched libraries. For that reason, you should always use those functions when adding support for compilers and linkers.

(Need a note about UNIX)

Custom generator classes

The standard generators allows you to specify source and target types, action, and a set of flags. If you need anything more complex, you need to create a new generator class with your own logic. Then, you have to create an instance of that class and register it. Here's an example how you can create your own generator class:

```
class custom-generator : generator
{
    rule __init__ ( * : * )
    {
        generator.__init__ $(1) : $(2) : $(3) : $(4) : $(5) : $(6) : $(7) : $(8) : $(9) ;
    }
}

generators.register
[ new custom-generator verbatim.inline-file : VERBATIM : CPP ] ;
```

This generator will work exactly like the `verbatim.inline-file` generator we've defined above, but it's possible to customize the behaviour by overriding methods of the generator class.

There are two methods of interest. The `run` method is responsible for the overall process - it takes a number of source targets, converts them the the right types, and creates the result. The `generated-targets` method is called when all sources are converted to the right types to actually create the result.

The `generated-target` method can be overridden when you want to add additional properties to the generated targets or use additional sources. For a real-life example, suppose you have a program analysis tool that should be given a name of executable and the list of all sources. Naturally, you don't want to list all source files manually. Here's how the `generated-target` method can find the list of sources automatically:

```
class itrace-generator : generator {
....
    rule generated-targets ( sources + : property-set : project name ? )
    {
        local leaves ;
        local temp = [ virtual-target.traverse $(sources[1]) : : include-sources ] ;
        for local t in $(temp)
        {
            if ! [ $(t).action ]
```

```

        {
            leaves += $(t) ;
        }
    }
    return [ generator.generated-targets $(sources) $(leafs)
            : $(property-set) : $(project) $(name) ] ;
}
}
generators.register [ new itrace-generator nm.itrace : EXE : ITRACE ] ;

```

The `generated-targets` method will be called with a single source target of type `EXE`. The call to `virtual-target.traverse` will return all targets the executable depends on, and we further find files that are not produced from anything. The found targets are added to the sources.

The `run` method can be overridden to completely customize the way generator works. In particular, the conversion of sources to the desired types can be completely customized. Here's another real example. Tests for the Boost Python library usually consist of two parts: a Python program and a C++ file. The C++ file is compiled to Python extension that is loaded by the Python program. But in the likely case that both files have the same name, the created Python extension must be renamed. Otherwise, Python program will import itself, not the extension. Here's how it can be done:

```

rule run ( project name ? : property-set : sources * : multiple ? )
{
    local python ;
    for local s in $(sources)
    {
        if [ $(s).type ] = PY
        {
            python = $(s) ;
        }
    }

    local libs ;
    for local s in $(sources)
    {
        if [ type.is-derived [ $(s).type ] LIB ]
        {
            libs += $(s) ;
        }
    }

    local new-sources ;
    for local s in $(sources)
    {
        if [ type.is-derived [ $(s).type ] CPP ]
        {
            local name = [ $(s).name ] ;    # get the target's basename
            if $(name) = [ $(python).name ]
            {
                name = $(name)_ext ;        # rename the target
            }
            new-sources += [ generators.construct $(project) $(name) :
                            PYTHON_EXTENSION : $(property-set) : $(s) $(libs) ] ;
        }
    }

    result = [ construct-result $(python) $(new-sources) : $(project) $(name)
              : $(property-set) ] ;
}

```

First, we separate all source into python files, libraries and C++ sources. For each C++ source we create a separate Python extension by calling `generators.construct` and passing the C++ source and the libraries. At this

point, we also change the extension's name, if necessary.

Features

Often, we need to control the options passed the invoked tools. This is done with features. Consider an example:

```
# Declare a new free feature
import feature : feature ;
feature verbatim-options : : free ;

# Cause the value of the 'verbatim-options' feature to be
# available as 'OPTIONS' variable inside verbatim.inline-file
import toolset : flags ;
flags verbatim.inline-file OPTIONS <verbatim-options> ;

# Use the "OPTIONS" variable
actions inline-file
{
    "./inline-file.py" $(OPTIONS) $(<) $(>)
}
```

We first define a new feature. Then, the `flags` invocation says that whenever `verbatim.inline-file` action is run, the value of the `verbatim-options` feature will be added to the `OPTIONS` variable, and can be used inside the action body. You'd need to consult online help (`--help`) to find all the features of the `toolset.flags` rule.

Although you can define any set of features and interpret their values in any way, `Boost.Build` suggests the following coding standard for designing features.

Most features should have a fixed set of values that is portable (tool neutral) across the class of tools they are designed to work with. The user does not have to adjust the values for a exact tool. For example, `<optimization>speed` has the same meaning for all C++ compilers and the user does not have to worry about the exact options passed to the compiler's command line.

Besides such portable features there are special 'raw' features that allow the user to pass any value to the command line parameters for a particular tool, if so desired. For example, the `<cxxflags>` feature allows you to pass any command line options to a C++ compiler. The `<include>` feature allows you to pass any string preceded by `-I` and the interpretation is tool-specific. (See the section called "Can I get output of external program as a variable in a Jamfile?" for an example of very smart usage of that feature). Of course one should always strive to use portable features, but these are still be provided as a backdoor just to make sure `Boost.Build` does not take away any control from the user.

Using portable features is a good idea because:

- When a portable feature is given a fixed set of values, you can build your project with two different settings of the feature and `Boost.Build` will automatically use two different directories for generated files. `Boost.Build` does not try to separate targets built with different raw options.
- Unlike with "raw" features, you don't need to use specific command-line flags in your Jamfile, and it will be more likely to work with other tools.

Steps for adding a feature

Adding a feature requires three steps:

1. Declaring a feature. For that, the "feature.feature" rule is used. You have to decide on the set of feature attrib-

utes:

- if a feature has several values and significantly affects the build, make it “propagated,” so that the whole project is built with the same value by default
 - if a feature does not have a fixed list of values, it must be “free.” For example, the `include` feature is a free feature.
 - if a feature is used to refer to a path relative to the Jamfile, it must be a “path” feature. `include` is also a path feature.
 - if feature is used to refer to some target, it must be a “dependency” feature.
2. Representing the feature value in a target-specific variable. Build actions are command templates modified by Boost.Jam variable expansions. The `toolset.flags` rule sets a target-specific variable to the value of a feature.
 3. Using the variable. The variable set in step 2 can be used in a build action to form command parameters or files.

Another example

Here's another example. Let's see how we can make a feature that refers to a target. For example, when linking dynamic libraries on windows, one sometimes needs to specify "DEF file", telling what functions should be exported. It would be nice to use this file like this:

```
lib a : a.cpp : <def-file>a.def ;
```

Actually, this feature is already supported, but anyway...

1. Since the feature refers to a target, it must be "dependency".

```
feature def-file : : free dependency ;
```

2. One of the toolsets that cares about DEF files is `msvc`. The following line should be added to it.

```
flags msvc.link DEF_FILE <def-file> ;
```

3. Since the `DEF_FILE` variable is not used by the `msvc.link` action, we need to modify it to be:

```
actions link bind DEF_FILE
{
    $(LD) .... /DEF:$(DEF_FILE) ....
}
```

Note the `bind DEF_FILE` part. It tells `bjam` to translate the internal target name in `DEF_FILE` to a corresponding filename in the `link` action. Without it the expansion of `$(DEF_FILE)` would be a strange symbol that is not likely to make sense for the linker.

We've almost done, but should stop for a small workaround. Add the following code to `msvc.jam`

```
rule link
{
    DEPENDS $(<) : [ on $(<) return $(DEF_FILE) ] ;
}
```

This is needed to accomodate some bug in bjam, which hopefully will be fixed one day.

Variants and composite features.

Sometimes you want to create a shortcut for some set of features. For example, `release` is a value of `<variant>` and is a shortcut for a set of features.

It is possible to define your own build variants. For example:

```
variant crazy : <optimization>speed <inlining>off
               <debug-symbols>on <profiling>on ;
```

will define a new variant with the specified set of properties. You can also extend an existing variant:

```
variant super_release : release : <define>USE_ASM ;
```

In this case, `super_release` will expand to all properties specified by `release`, and the additional one you've specified.

You are not restricted to using the `variant` feature only. Here's example that defines a brand new feature:

```
feature parallelism : mpi fake none : composite link-incompatible ;
feature.compose <parallelism>mpi : <library>/mpi//mpi/<parallelism>none ;
feature.compose <parallelism>fake : <library>/mpi//fake/<parallelism>none ;
```

This will allow you to specify value of feature `parallelism`, which will expand to link to the necessary library.

Main target rules

A main target rule (e.g “`exe`” Or “`lib`”) creates a top-level target. It's quite likely that you'll want to declare your own and there are two ways to do that.

The first way applies when your target rule should just produce a target of specific type. In that case, a rule is already defined for you! When you define a new type, `Boost.Build` automatically defines a corresponding rule. The name of the rule is obtained from the name of the type, by downcasing all letters and replacing underscores with dashes. For example, if you create a module `obfuscate.jam` containing:

```
import type ;
type.register OBFUSCATED_CPP : ocpp ;

import generators ;
generators.register-standard obfuscate.file : CPP : OBFUSCATED_CPP ;
```

and import that module, you'll be able to use the rule "obfuscated-cpp" in Jamfiles, which will convert source to the `OBFUSCATED_CPP` type.

The second way is to write a wrapper rule that calls any of the existing rules. For example, suppose you have only

one library per directory and want all cpp files in the directory to be compiled into that library. You can achieve this effect with:

```
lib codegen : [ glob *.cpp ] ;
```

but if you want to make it even simpler, you could add the following definition to the `project-root.jam` file:

```
rule glib ( name : extra-sources * : requirements * )
{
    lib $(name) : [ glob *.cpp ] $(extra-sources) : $(requirements) ;
}
```

which would allow you to reduce the Jamfile to

```
glib codegen ;
```

Note that because you can associate a custom generator with a target type, the logic of building can be rather complex. For example, the `boostbook` module declares a target type `BOOSTBOOK_MAIN` and a custom generator for that type. You can use that as example if your main target rule is non-trivial.

Toolset modules

If your extensions will be used only on one project, they can be placed in a separate `.jam` file that will be imported by your `project-root.jam`. If the extensions will be used on many projects, users will thank you for a finishing touch.

The `using` rule provides a standard mechanism for loading and configuring extensions. To make it work, your module should provide an `init` rule. The rule will be called with the same parameters that were passed to the `using` rule. The set of allowed parameters is determined by you. For example, you can allow the user to specify paths, tool versions, and other options.

Here are some guidelines that help to make Boost.Build more consistent:

- The `init` rule should never fail. Even if the user provided an incorrect path, you should emit a warning and go on. Configuration may be shared between different machines, and wrong values on one machine can be OK on another.
- Prefer specifying the command to be executed to specifying the tool's installation path. First of all, this gives more control: it's possible to specify

```
/usr/bin/g++-snapshot
time g++
```

as the command. Second, while some tools have a logical "installation root", it's better if user doesn't have to remember whether a specific tool requires a full command or a path.

- Check for multiple initialization. A user can try to initialize the module several times. You need to check for this and decide what to do. Typically, unless you support several versions of a tool, duplicate initialization is a user error. If the tool's version can be specified during initialization, make sure the version is either always specified, or never specified (in which case the tool is initialized only once). For example, if you allow:

```
using yfc ;
```

```
using yfc : 3.3 ;  
using yfc : 3.4 ;
```

Then it's not clear if the first initialization corresponds to version 3.3 of the tool, version 3.4 of the tool, or some other version. This can lead to building twice with the same version.

- If possible, `init` must be callable with no parameters. In which case, it should try to autodetect all the necessary information, for example, by looking for a tool in `PATH` or in common installation locations. Often this is possible and allows the user to simply write:

```
using yfc ;
```

- Consider using facilities in the `tools/common` module. You can take a look at how `tools/gcc.jam` uses that module in the `init` rule.

Detailed reference

General information

Initialization

bjam's first job upon startup is to load the Jam code that implements the build system. To do this, it searches for a file called `boost-build.jam`, first in the invocation directory, then in its parent and so forth up to the filesystem root, and finally in the directories specified by the environment variable `BOOST_BUILD_PATH`. When found, the file is interpreted, and should specify the build system location by calling the `boost-build` rule:

```
rule boost-build ( location ? )
```

If `location` is a relative path, it is treated as relative to the directory of `boost-build.jam`. The directory specified by that `location` and the directories in `BOOST_BUILD_PATH` are then searched for a file called `bootstrap.jam`, which is expected to bootstrap the build system. This arrangement allows the build system to work without any command-line or environment variable settings. For example, if the build system files were located in a directory "build-system/" at your project root, you might place a `boost-build.jam` at the project root containing:

```
boost-build build-system ;
```

In this case, running `bjam` anywhere in the project tree will automatically find the build system.

The default `bootstrap.jam`, after loading some standard definitions, loads two files, which can be provided/customised by user: `site-config.jam` and `user-config.jam`.

Locations where those files a search are summarized below:

Table 1. Search paths for configuration files

	<code>site-config.jam</code>	<code>user-config.jam</code>
Linux	/etc \$HOME \$BOOST_BUILD_PATH	\$HOME \$BOOST_BUILD_PATH
Windows	%SystemRoot% %HOMEDRIVE%%HOMEPATH% %HOME% %BOOST_BUILD_PATH%	%HOMEDRIVE%%HOMEPATH% %HOME% %BOOST_BUILD_PATH%

Boost.Build comes with default versions of those files, which can serve as templates for customized versions.

Command line

The command line may contain:

- Jam options,
- Boost.Build options,
- Command line arguments

Command line arguments

Command line arguments specify targets and build request using the following rules.

- An argument that does not contain slashes or the = symbol is either a value of an implicit feature or of a target to be built. It is taken to be value of a feature if an appropriate feature exists. Otherwise, it is considered a target id. Building the special target name “clean” has the same effect as using the `--clean` option.
- An argument containing either slashes or the = symbol specifies a number of build request elements (see ???). In its simplest form, it's just a set of properties, separated by slashes, which become a single build request element, for example:

```
borland/<runtime-link>static
```

A more complex form can be used to save typing. For example, instead of

```
borland/runtime-link=static borland/runtime-link=dynamic
```

one can use

```
borland/runtime-link=static,dynamic
```

Exactly, the conversion from argument to build request elements is performed by (1) splitting the argument at each slash, (2) converting each split part into a set of properties and (3) taking all possible combinations of the property sets. Each split part should have the either the form

```
feature-name=feature-value1[" , "feature-valueN]*
```

or, in case of implicit features

```
feature-value1[" , "feature-valueN; ]*
```

will be converted into the property set

```
<feature-name>feature-value1 .... <feature-name>feature-valueN
```

For example, the command line

```
target1 debug gcc/runtime-link=dynamic,static
```

would cause target called `target1` to be rebuilt in debug mode, except that for `gcc`, both dynamically and static-

ally linked binaries would be created.

Command line options

All of the Boost.Build options start with the "--" prefix. They are described in the following table.

Table 2. Command line options

Option	Description
<code>--version</code>	Prints information on Boost.Build and Boost.Jam versions.
<code>--help</code>	Access to the online help system. This prints general information on how to use the help system with additional <code>--help*</code> options.
<code>--clean</code>	Removes everything instead of building. Unlike <code>clean</code> target in <code>make</code> , it is possible to clean only some targets.
<code>--debug</code>	Enables internal checks.
<code>--dump-projects</code>	Cause the project structure to be output.
<code>--no-error-backtrace</code>	Don't print backtrace on errors. Primary useful for testing.
<code>--ignore-config</code>	Do not load <code>site-config.jam</code> and <code>user-config.jam</code>

Writing Jamfiles

This section describes specific information about writing Jamfiles.

Generated headers

Usually, Boost.Build handles implicit dependencies completely automatically. For example, for C++ files, all `#include` statements are found and handled. The only aspect where user help might be needed is implicit dependency on generated files.

By default, Boost.Build handles such dependencies within one main target. For example, assume that main target "app" has two sources, "app.cpp" and "parser.y". The latter source is converted into "parser.c" and "parser.h". Then, if "app.cpp" includes "parser.h", Boost.Build will detect this dependency. Moreover, since "parser.h" will be generated into a build directory, the path to that directory will automatically added to include path.

Making this mechanism work across main target boundaries is possible, but imposes certain overhead. For that reason, if there's implicit dependency on files from other main targets, the `<implicit-dependency>` [link] feature must be used, for example:

```
lib parser : parser.y ;
exe app : app.cpp : <implicit-dependency>parser ;
```

The above example tells the build system that when scanning all sources of "app" for implicit-dependencies, it should consider targets from "parser" as potential dependencies.

Build process

The general overview of the build process was given in the user documentation. This section provides additional details, and some specific rules.

To recap, building a target with specific properties includes the following steps:

1. applying default build,
2. selecting the main target alternative to use,
3. determining "common" properties
4. building targets referred by the sources list and dependency properties
5. adding the usage requirements produces when building dependencies to the "common" properties
6. building the target using generators
7. computing the usage requirements to be returned

Alternative selection

When there are several alternatives, one of them must be selected. The process is as follows:

1. For each alternative *condition* is defined as the set of base properties in requirements. [Note: it might be better to specify the condition explicitly, as in conditional requirements].
2. An alternative is viable only if all properties in condition are present in build request.
3. If there's one viable alternative, it's chosen. Otherwise, an attempt is made to find one best alternative. An alternative a is better than another alternative b, iff set of properties in b's condition is strict subset of the set of properties of a's condition. If there's one viable alternative, which is better than all other, it's selected. Otherwise, an error is reported.

Determining common properties

The "common" properties is a somewhat artificial term. Those are the intermediate property set from which both the build request for dependencies and properties for building the target are derived.

Since default build and alternatives are already handled, we have only two inputs: build requests and requirements. Here are the rules about common properties.

1. Non-free feature can have only one value
2. A non-conditional property in requirement is always present in common properties.
3. A property in build request is present in common properties, unless (2) tells otherwise.
4. If either build request, or requirements (non-conditional or conditional) include an expandable property (either composite, or property with specified subfeature value), the behaviour is equivalent to explicitly adding all expanded properties to build request or requirements.
5. If requirements include a conditional property, and condition of this property is true in context of common properties, then the conditional property should be in common properties as well.

6. If no value for a feature is given by other rules here, it has default value in common properties.

Those rules are declarative, they don't specify how to compute the common properties. However, they provide enough information for the user. The important point is the handling of conditional requirements. The condition can be satisfied either by property in build request, by non-conditional requirements, or even by another conditional property. For example, the following example works as expected:

```
exe a : a.cpp
      : <toolset>gcc:<variant>release
      <variant>release:<define>FOO ;
```

Definitions

Features and properties

A *feature* is a normalized (toolset-independent) aspect of a build configuration, such as whether inlining is enabled. Feature names may not contain the '>' character.

Each feature in a build configuration has one or more associated *values*. Feature values for non-free features may not contain the '<', ':', or '=' characters. Feature values for free features may not contain the '<' character.

A *property* is a (feature,value) pair, expressed as <feature>value.

A *subfeature* is a feature that only exists in the presence of its parent feature, and whose identity can be derived (in the context of its parent) from its value. A subfeature's parent can never be another subfeature. Thus, features and their subfeatures form a two-level hierarchy.

A *value-string* for a feature **F** is a string of the form value-subvalue1-subvalue2...-subvalueN, where value is a legal value for **F** and subvalue1...subvalueN are legal values of some of **F**'s subfeatures. For example, the properties <toolset>gcc <toolset-version>3.0.1 can be expressed more concisely using a value-string, as <toolset>gcc-3.0.1.

A *property set* is a set of properties (i.e. a collection without duplicates), for instance: <toolset>gcc <runtime-link>static.

A *property path* is a property set whose elements have been joined into a single string separated by slashes. A property path representation of the previous example would be <toolset>gcc/<runtime-link>static.

A *build specification* is a property set that fully describes the set of features used to build a target.

Property Validity

For free features, all values are valid. For all other features, the valid values are explicitly specified, and the build system will report an error for the use of an invalid feature-value. Subproperty validity may be restricted so that certain values are valid only in the presence of certain other subproperties. For example, it is possible to specify that the <gcc-target>mingw property is only valid in the presence of <gcc-version>2.95.2.

Feature Attributes

Each feature has a collection of zero or more of the following attributes. Feature attributes are low-level descriptions of how the build system should interpret a feature's values when they appear in a build request. We also refer to the attributes of properties, so that an *incidental* property, for example, is one whose feature has the *incidental* attribute.

- *incidental*

Incidental features are assumed not to affect build products at all. As a consequence, the build system may use the same file for targets whose build specification differs only in incidental features. A feature that controls a compiler's warning level is one example of a likely incidental feature.

Non-incidental features are assumed to affect build products, so the files for targets whose build specification differs in non-incidental features are placed in different directories as described in "target paths" below. [where?]

- *propagated*

Features of this kind are propagated to dependencies. That is, if a main target is built using a propagated property, the build systems attempts to use the same property when building any of its dependencies as part of that main target. For instance, when an optimized executable is requested, one usually wants it to be linked with optimized libraries. Thus, the `<optimization>` feature is propagated.

- *free*

Most features have a finite set of allowed values, and can only take on a single value from that set in a given build specification. Free features, on the other hand, can have several values at a time and each value can be an arbitrary string. For example, it is possible to have several preprocessor symbols defined simultaneously:

```
<define>NDEBUG=1 <define>HAS_CONFIG_H=1
```

- *optional*

An optional feature is a feature that is not required to appear in a build specification. Every non-optional non-free feature has a default value that is used when a value for the feature is not otherwise specified, either in a target's requirements or in the user's build request. [A feature's default value is given by the first value listed in the feature's declaration. -- move this elsewhere - dwa]

- *symmetric*

A symmetric feature's default value is not automatically included in build variants. Normally a feature only generates a subvariant directory when its value differs from the value specified by the build variant, leading to an asymmetric subvariant directory structure for certain values of the feature. A symmetric feature, when relevant to the toolset, always generates a corresponding subvariant directory.

- *path*

The value of a path feature specifies a path. The path is treated as relative to the directory of Jamfile where path feature is used and is translated appropriately by the build system when the build is invoked from a different directory

- *implicit*

Values of implicit features alone identify the feature. For example, a user is not required to write "`<toolset>gcc`", but can simply write "gcc". Implicit feature names also don't appear in variant paths, although the values do. Thus: `bin/gcc/...` as opposed to `bin/toolset-gcc/...` There should typically be only a few such features, to avoid possible name clashes.

- *composite*

Composite features actually correspond to groups of properties. For example, a build variant is a composite feature. When generating targets from a set of build properties, composite features are recursively expanded and *added* to the build property set, so rules can find them if necessary. Non-composite non-free features override com-

ponents of composite features in a build property set.

- *dependency*

The value of dependency feature if a target reference. When used for building of a main target, the value of dependency feature is treated as additional dependency.

For example, dependency features allow to state that library A depends on library B. As the result, whenever an application will link to A, it will also link to B. Specifying B as dependency of A is different from adding B to the sources of A.

Features that are neither free nor incidental are called *base* features.

Feature Declaration

The low-level feature declaration interface is the `feature` rule from the `feature` module:

```
rule feature ( name : allowed-values * : attributes * )
```

A feature's allowed-values may be extended with the `feature.extend` rule.

Build Variants

A build variant, or (simply variant) is a special kind of composite feature that automatically incorporates the default values of features that . Typically you'll want at least two separate variants: one for debugging, and one for your release code. [Volodya says: "Yea, we'd need to mention that it's a composite feature and describe how they are declared, in particular that default values of non-optional features are incorporated into build variant automatically. Also, do we want some variant inheritance/extension/templates. I don't remember how it works in V1, so can't document this for V2.". Will clean up soon -DWA]

Property refinement

When a target with certain properties is requested, and that target requires some set of properties, it is needed to find the set of properties to use for building. This process is called *property refinement* and is performed by these rules

1. Each property in the required set is added to the original property set
2. If the original property set includes property with a different value of non free feature, that property is removed.

Conditional properties

Sometime it's desirable to apply certain requirements only for a specific combination of other properties. For example, one of compilers that you use issues a pointless warning that you want to suppress by passing a command line option to it. You would not want to pass that option to other compilers. Conditional properties allow you to do just that. Their syntax is:

```
property ( "," property ) * ":" property
```

For example, the problem above would be solved by:

```
exe hello : hello.cpp : <toolset>yfc:<cxxflags>-disable-pointless-warning ;
```

The syntax also allows several properties in the condition, for example:

```
exe hello : hello.cpp : <os>NT,<toolset>gcc:<link>static ;
```

Target identifiers and references

Target identifier is used to denote a target. The syntax is:

```
target-id -> (project-id | target-name | file-name )
              | (project-id | directory-name) "/" target-name
project-id -> path
target-name -> path
file-name -> path
directory-name -> path
```

This grammar allows some elements to be recognized as either

- project id (at this point, all project ids start with slash).
- name of target declared in current Jamfile (note that target names may include slash).
- a regular file, denoted by absolute name or name relative to project's sources location.

To determine the real meaning a check is made if project-id by the specified name exists, and then if main target of that name exists. For example, valid target ids might be:

```
a                -- target in current project
lib/b.cpp        -- regular file
/boost/thread    -- project "/boost/thread"
/home/ghost/build/lr_library//parser -- target in specific project
```

Rationale: Target is separated from project by special separator (not just slash), because:

- It emphasises that projects and targets are different things.
- It allows to have main target names with slashes.

Target reference is used to specify a source target, and may additionally specify desired properties for that target. It has this syntax:

```
target-reference -> target-id [ "/" requested-properties ]
requested-properties -> property-path
```

For example,

```
exe compiler : compiler.cpp libs/cmdline/<optimization>space ;
```

would cause the version of `cmdline` library, optimized for space, to be linked in even if the `compiler` executable is build with optimization for speed.

Generators

Warning

The information in this section is likely to be outdated and misleading.

To construct a main target with given properties from sources, it is required to create a dependency graph for that main target, which will also include actions to be run. The algorithm for creating the dependency graph is described here.

The fundamental concept is *generator*. It encapsulates the notion of build tool and is capable to converting a set of input targets into a set of output targets, with some properties. Generator matches a build tool as closely as possible: it works only when the tool can work with requested properties (for example, `msvc` compiler can't work when requested toolset is `gcc`), and should produce exactly the same targets as the tool (for example, if Borland's linker produces additional files with debug information, generator should also).

Given a set of generators, the fundamental operation is to construct a target of a given type, with given properties, from a set of targets. That operation is performed by rule `generators.construct` and the used algorithm is described below.

Selecting and ranking viable generators

Each generator, in addition to target types that it can produce, have attribute that affects its applicability in particular situation. Those attributes are:

1. Required properties, which are properties absolutely necessary for the generator to work. For example, generator encapsulating the `gcc` compiler would have `<toolset>gcc` as required property.
2. Optional properties, which increase the generators suitability for a particular build.

Generator's required and optional properties may not include either free or incidental properties. (Allowing this would greatly complicate caching targets).

When trying to construct a target, the first step is to select all possible generators for the requested target type, which required properties are a subset of requested properties. Generators that were already selected up the call stack are excluded. In addition, if any composing generators were selected up the call stack, all other composing generators are ignored (TODO: define composing generators). The found generators are assigned a rank, which is the number of optional properties present in requested properties. Finally, generators with highest rank are selected for further processing.

Running generators

When generators are selected, each is run to produce a list of created targets. This list might include targets that are not of requested types, because generators create the same targets as some tool, and tool's behaviour is fixed. (Note: should specify that in some cases we actually want extra targets). If generator fails, it returns an empty list. Generator is free to call 'construct' again, to convert sources to the types it can handle. It also can pass modified properties to 'construct'. However, a generator is not allowed to modify any propagated properties, otherwise when actually consuming properties we might discover that the set of propagated properties is different from what was used for building sources.

For all targets that are not of requested types, we try to convert them to requested type, using a second call to `construct`. This is done in order to support transformation sequences where single source file expands to several later. See this message for details.

Selecting dependency graph

After all generators are run, it is necessary to decide which of successful invocation will be taken as final result. At the moment, this is not done. Instead, it is checked whether all successful generator invocation returned the same target list. Error is issued otherwise.

Property adjustment

Because target location is determined by the build system, it is sometimes necessary to adjust properties, in order to not break actions. For example, if there's an action that generates a header, say `"a_parser.h"`, and a source file `"a.cpp"` which includes that file, we must make everything work as if `a_parser.h` is generated in the same directory where it would be generated without any subvariants.

Correct property adjustment can be done only after all targets are created, so the approach taken is:

1. When dependency graph is constructed, each action can be assigned a rule for property adjustment.
2. When virtual target is actualized, that rule is run and return the final set of properties. At this stage it can use information of all created virtual targets.

In case of quoted includes, no adjustment can give 100% correct results. If target dirs are not changed by build system, quoted includes are searched in `"."` and then in include path, while angle includes are searched only in include path. When target dirs are changed, we'd want to make quoted includes to be search in `"."` then in additional dirs and then in the include path and make angle includes be searched in include path, probably with additional paths added at some position. Unless, include path already has `"."` as the first element, this is not possible. So, either generated headers should not be included with quotes, or first element of include path should be `"."`, which essentially erases the difference between quoted and angle includes. **Note:** the only way to get `"."` as include path into compiler command line is via verbatim compiler option. In all other case, `Boost.Build` will convert `"."` into directory where it occurs.

Transformations cache

Under certain conditions, an attempt is made to cache results of transformation search. First, the sources are replaced with targets with special name and the found target list is stored. Later, when properties, requested type, and source type are the same, the store target list is retrieved and cloned, with appropriate change in names.

Frequently Asked Questions

I'm getting "Duplicate name of actual target" error. What does it mean?

The most likely case is that you're trying to compile the same file twice, with almost the same, but differing properties. For example:

```
exe a : a.cpp : <include>/usr/local/include ;
exe b : a.cpp ;
```

The above snippet requires two different compilations of 'a.cpp', which differ only in 'include' property. Since the 'include' property is free, Boost.Build can't generate two objects files into different directories. On the other hand, it's dangerous to compile the file only once -- maybe you really want to compile with different includes.

To solve this issue, you need to decide if file should be compiled once or twice.

1. Two compile file only once, make sure that properties are the same:

```
exe a : a.cpp : <include>/usr/local/include ;
exe b : a.cpp : <include>/usr/local/include ;
```

2. If changing the properties is not desirable, for example if 'a' and 'b' target have other sources which need specific properties, separate 'a.cpp' into it's own target:

```
obj a_obj : a.cpp : <include>/usr/local/include ;
exe a : a_obj ;
```

3. To compile file twice, you can make the object file local to the main target:

```
exe a : [ obj a_obj : a.cpp ] : <include>/usr/local/include ;
exe b : [ obj a_obj : a.cpp ] ;
```

A good question is why Boost.Build can't use some of the above approaches automatically. The problem is that such magic would require additional implementation complexities and would only help in half of the cases, while in other half we'd be silently doing the wrong thing. It's simpler and safe to ask user to clarify his intention in such cases.

Accessing environment variables

Many users would like to use environment variables in Jamfiles, for example, to control location of external libraries. In many cases you better declare those external libraries in the site-config.jam file, as documented in the recipes section. However, if the users already have the environment variables set up, it's not convenient to ask them to set up site-config.jam files as well, and using environment variables might be reasonable.

In Boost.Build V2, each Jamfile is a separate namespace, and the variables defined in environment is imported into the global namespace. Therefore, to access environment variable from Jamfile, you'd need the following code:

```
import modules ;
local SOME_LIBRARY_PATH = [ modules.peek : SOME_LIBRARY_PATH ] ;
exe a : a.cpp : <include>$(SOME_LIBRARY_PATH) ;
```

How to control properties order?

For internal reasons, Boost.Build sorts all the properties alphabetically. This means that if you write:

```
exe a : a.cpp : <include>b <include>a ;
```

then the command line will first mention the "a" include directory, and then "b", even though they are specified in the opposite order. In most cases, the user doesn't care. But sometimes the order of includes, or other properties, is important. For example, if one uses both the C++ Boost library and the "boost-sandbox" (libraries in development), then include path for boost-sandbox must come first, because some headers may override ones in C++ Boost. For such cases, a special syntax is provided:

```
exe a : a.cpp : <include>a&&b ;
```

The && symbols separate values of an property, and specify that the order of the values should be preserved. You are advised to use this feature only when the order of properties really matters, and not as a convenient shortcut. Using it everywhere might negatively affect performance.

How to control the library order on Unix?

On the Unix-like operating systems, the order in which static libraries are specified when invoking the linker is important, because by default, the linker uses one pass through the libraries list. Passing the libraries in the incorrect order will lead to a link error. Further, this behaviour is often used to make one library override symbols from another. So, sometimes it's necessary to force specific order of libraries.

Boost.Build tries to automatically compute the right order. The primary rule is that if library a "uses" library b, then library a will appear on the command line before library b. Library a is considered to use b if b is present either in the sources of a or in its requirements. To explicitly specify the use relationship one can use the <use> feature. For example, both of the following lines will cause a to appear before b on the command line:

```
lib a : a.cpp b ;
lib a : a.cpp : <use>b ;
```

The same approach works for searched libraries, too:

```
lib z ;
lib png : : <use>z ;
exe viewer : viewer png z ;
```

Can I get output of external program as a variable in a Jamfile?

From time to time users ask how to run an external program and save the result in Jamfile variable, something like:


```
local gtk_includes = [ RUN_COMMAND gtk-config ] ;
```

Unfortunately, this is not possible at the moment. However, if the result of command invocation is to be used in a command to some tool, and you're working on Unix, the following workaround is possible.

```
alias gtk+-2.0 : : : :
    <cflags>"`pkg-config --cflags gtk+-2.0`"
    <inkflags>"`pkg-config --libs gtk+-2.0`"
;
```

If you use the "gtk+-2.0" target in sources, then the properties specified above will be added to the build properties and eventually will appear in the command line. Unix command line shell processes the backticks quoting by running the tool and using its output -- which is what's desired in that case. Thanks to Daniel James for sharing this approach.

How to get the project-root location?

You might want to use the location of the project-root in your Jamfiles. To do it, you'd need to declare path constant in your project-root.jam:

```
path-constant TOP : . ;
```

After that, the TOP variable can be used in every Jamfile.

How to change compilation flags for one file?

If one file must be compiled with special options, you need to explicitly declare an `obj` target for that file and then use that target in your `exe` or `lib` target:

```
exe a : a.cpp b ;
obj b : b.cpp : <optimization>off ;
```

Of course you can use other properties, for example to specify specific compiler options:

```
exe a : a.cpp b ;
obj b : b.cpp : <cflags>-g ;
```

You can also use conditional properties for finer control:

```
exe a : a.cpp b ;
obj b : b.cpp : <variant>release:<optimization>off ;
```

Why are the `dll-path` and `hardcode-dll-paths` properties useful?

(This entry is specific to Unix system.) Before answering the questions, let's recall a few points about shared libraries. Shared libraries can be used by several applications, or other libraries, without physically including the library in the application. This can greatly decrease the total size of applications. It's also possible to upgrade a shared library when the application is already installed. Finally, shared linking can be faster.

However, the shared library must be found when the application is started. The dynamic linker will search in a system-defined list of paths, load the library and resolve the symbols. Which means that you should either change the system-defined list, given by the `LD_LIBRARY_PATH` environment variable, or install the libraries to a system location. This can be inconvenient when developing, since the libraries are not yet ready to be installed, and cluttering system paths is undesirable. Luckily, on Unix there's another way.

An executable can include a list of additional library paths, which will be searched before system paths. This is excellent for development, because the build system knows the paths to all libraries and can include them in executables. That's done when the `hardcode-dll-paths` feature has the `true` value, which is the default. When the executables should be installed, the story is different.

Obviously, installed executable should not hardcode paths to your development tree. (The `stage` rule explicitly disables the `hardcode-dll-paths` feature for that reason.) However, you can use the `dll-path` feature to add explicit paths manually. For example:

```
stage installed : application : <dll-path>/usr/lib/snake
                                <location>/usr/bin ;
```

will allow the application to find libraries placed to `/usr/lib/snake`.

If you install libraries to a nonstandard location and add an explicit path, you get more control over libraries which will be used. A library of the same name in a system location will not be inadvertently used. If you install libraries to a system location and do not add any paths, the system administrator will have more control. Each library can be individually upgraded, and all applications will use the new library.

Which approach is best depends on your situation. If the libraries are relatively standalone and can be used by third party applications, they should be installed in the system location. If you have lots of libraries which can be used only by your application, it makes sense to install it to a nonstandard directory and add an explicit path, like the example above shows. Please also note that guidelines for different systems differ in this respect. The Debian guidelines prohibit any additional search paths, and Solaris guidelines suggest that they should always be used.

Targets in `site-config.jam`

It is desirable to declare standard libraries available on a given system. Putting target declaration in Jamfile is not really good, since locations of the libraries can vary. The solution is to put the following to `site-config.jam`.

```
import project ;
project.initialize $(__name__) ;
project site-config ;
lib zlib : : <name>z ;
```

The second line allows this module to act as project. The third line gives id to this project — it really has no location and cannot be used otherwise. The fourth line just declares a target. Now, one can write:

```
exe hello : hello.cpp /site-config//zlib ;
```

in any Jamfile.

Appendix A. Boost.Build v2 architecture

This document is work-in progress. Don't expect much from it yet.

Overview

The Boost.Build code is structured in four different components: "kernel", "util", "build" and "tools". The first two are relatively uninteresting, so we'll focus on the remaining pair. The "build" component provides classes necessary to declare targets, determine which properties should be used for their building, and for creating the dependency graph. The "tools" component provides user-visible functionality. It mostly allows to declare specific kind of main targets, and declare available tools, which are then used when creating the dependency graph.

The build layer

The build layer has just four main parts -- metatargets (abstract targets), virtual targets, generators and properties.

- Metatargets (see the "targets.jam" module) represent all the user-defined entities which can be built. The "meta" prefix signify that they don't really correspond to files -- depending of build request, they can produce different set of files. Metatargets are created when Jamfiles are loaded. Each metatarget has a `generate` method which is given a property set and produces virtual targets for the passed properties.
- Virtual targets (see the "virtual-targets.jam" module) correspond to the atomic things which can be updated -- most typically files.
- Properties are just (name, value) pairs, specified by the user and describing how the targets should be built. Properties are stored using the `property-set` class.
- Generators are the objects which encapsulate tools -- they can take a list of source virtual targets and produce new virtual targets from them.

The build process includes those steps:

1. Top-level code calls the `generate` method of a metatarget with some properties.
2. The metatarget combines the requested properties with requirements and passes the result, together with the list of sources, to the `generators.construct` function
3. A generator appropriate for the build properties is selected and its `run` method is called. The method returns a list of virtual targets
4. The targets are returned to the top level code. They are converted into bjam targets (via `virtual-target.actualize`) and passed to bjam for building.

Metatargets

There are several classes derived from "abstract-target". The "main-target" class represents top-level main target, the "project-target" acts like container for all main targets, and "basic-target" class is a base class for all further target types.

Since each main target can have several alternatives, all top-level target objects are just containers, referring to "real" main target classes. The type of that container is "main-target". For example, given:

```
alias a ;  
lib a : a.cpp : <toolset>gcc ;
```

we would have one top-level instance of "main-target-class", which will contain one instance of "alias-target-class" and one instance of "lib-target-class". The "generate" method of "main-target" decides which of the alternatives should be used, and call "generate" on the corresponding instance.

Each alternative is an instance of a class derived from "basic-target". The "basic-target.generate" does several things that are always should be done:

- Determines what properties should be used for building the target. This includes looking at requested properties, requirements, and usage requirements of all sources.
- Builds all sources
- Computes the usage requirements which should be passed back.

For the real work of constructing virtual target, a new method "construct" is called.

The "construct" method can be implemented in any way by classes derived from "basic-target", but one specific derived class plays the central role -- "typed-target". That class holds the desired type of file to be produced, and calls the generators modules to do the job.

This means that a specific metatarget subclass may avoid using generators at all. However, this is deprecated and we're trying to eliminate all such subclasses at the moment.

Note that the `build/targets.jam` file contains an UML diagram which might help.

Virtual targets

Virtual targets correspond to the atomic things which can be updated. Each virtual target can be assigned an updating action -- instance of the `action` class. The `action` class, in turn, contains a list of source targets, properties, and a name of bjam action block which should be executed.

We try hard to never create equal instances of the `virtual-target` class. Each code which creates virtual targets passes them through the `virtual-target.register` function, which detects if a target with the same name, sources, and properties was created. In that case, existing target is returned.

When all virtual targets are produced, they are "actualized". This means that the real file names are computed, and the commands that should be run are generated. This is done by the `virtual-target.actualize` method and the `action.actualize` methods. The first is conceptually simple, while the second needs additional explanation. The commands in bjam are generated in two-stage process. First, a rule with the appropriate name (for example "gcc.compile") is called and is given the names of targets. The rule sets some variables, like "OPTIONS". After that, the command string is taken, and variables are substituted, so use of OPTIONS inside the command string becomes the real compile options.

Boost.Build added a third stage to simplify things. It's now possible to automatically convert properties to appropriate assignments to variables. For example, `<debug-symbols>on` would add "-g" to the OPTIONS variable, without requiring to manually add this logic to `gcc.compile`. This functionality is part of the "toolset" module.

Note that the `build/virtual-targets.jam` file contains an UML diagram which might help.

Above, we noted that metatargets are built with a set of properties. That set is represented with the `property-set` class. An important point is that handling of property sets can get very expensive. For that reason, we make sure that for each set of (name, value) pairs only one `property-set` instance is created. The `property-set` uses extensive caching for all operation, so most work is avoided. The `property-set.create` is the factory function which should be used to create instances of the `property-set` class.

The tools layer

Write me!

Targets

NOTE: THIS SECTION IS NOT EXPECTED TO BE READ! There are two user-visible kinds of targets in Boost.Build. First are "abstract" — they correspond to things declared by user, for example, projects and executable files. The primary thing about abstract target is that it's possible to request them to be build with a particular values of some properties. Each combination of properties may possible yield different set of real file, so abstract target do not have a direct correspondence with files.

File targets, on the contrary, are associated with concrete files. Dependency graphs for abstract targets with specific properties are constructed from file targets. User has no way to create file targets, however it can specify rules that detect file type for sources, and also rules for transforming between file targets of different types. That information is used in constructing dependency graph, as described in the "next section". [link?] **Note:**File targets are not the same as targets in Jam sense; the latter are created from file targets at the latest possible moment. **Note:**"File target" is a proposed name for what we call virtual targets. It is more understandable by users, but has one problem: virtual targets can potentially be "phony", and not correspond to any file.

Dependency scanning

Dependency scanning is the process of finding implicit dependencies, like `#include` statements in C++. The requirements for right dependency scanning mechanism are:

- Support for different scanning algorithms. C++ and XML have quite different syntax for includes and rules for looking up included files.
- Ability to scan the same file several times. For example, single C++ file can be compiled with different include paths.
- Proper detection of dependencies on generated files.
- Proper detection of dependencies from generated file.

Support for different scanning algorithms

Different scanning algorithm are encapsulated by objects called "scanners". Please see the documentation for "scanner" module for more details.

Ability to scan the same file several times

As said above, it's possible to compile a C++ file twice, with different include paths. Therefore, include dependencies for those compilations can be different. The problem is that bjam does not allow several scans of the same target.

The solution in Boost.Build is straightforward. When a virtual target is converted to bjam target (via `virtual-target.actualize` method), we specify the scanner object to be used. The actualize method will create different bjam targets for different scanners.

All targets with specific scanner are made dependent on target without scanner, which target is always created. This is done in case the target is updated. The updating action will be associated with target without scanner, but if sources for that action are touched, all targets — with scanner and without should be considered outdated.

For example, assume that "a.cpp" is compiled by two compilers with different include path. It's also copied into some install location. In turn, it's produced from "a.verbatim". The dependency graph will look like:

```

a.o (<toolset>gcc)  <--(compile)-- a.cpp (scanner1)  ----+
a.o (<toolset>msvc) <--(compile)-- a.cpp (scanner2) ----|
a.cpp (installed copy) <--(copy) ----- a.cpp (no scanner)
                                                    ^
                                                    |
a.verbose -----+

```

Proper detection of dependencies on generated files.

This requirement breaks down to the following ones.

1. If when compiling "a.cpp" there's include of "a.h", the "dir" directory is in include path, and a target called "a.h" will be generated to "dir", then bjam should discover the include, and create "a.h" before compiling "a.cpp".
2. Since almost always Boost.Build generates targets to a "bin" directory, it should be supported as well. I.e. in the scenario above, Jamfile in "dir" might create a main target, which generates "a.h". The file will be generated to "dir/bin" directory, but we still have to recognize the dependency.

The first requirement means that when determining what "a.h" means, when found in "a.cpp", we have to iterate over all directories in include paths, checking for each one:

1. If there's file "a.h" in that directory, or
2. If there's a target called "a.h", which will be generated to that directory.

Classic Jam has built-in facilities for point (1) above, but that's not enough. It's hard to implement the right semantic without builtin support. For example, we could try to check if there's target called "a.h" somewhere in dependency graph, and add a dependency to it. The problem is that without search in include path, the semantic may be incorrect. For example, one can have an action which generated some "dummy" header, for system which don't have the native one. Naturally, we don't want to depend on that generated header on platforms where native one is included.

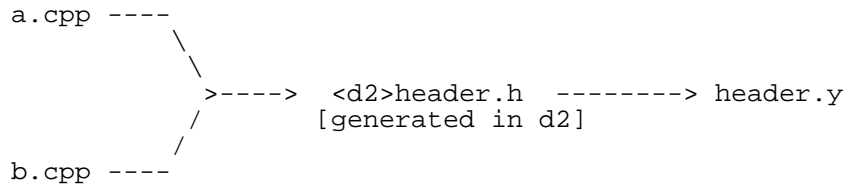
There are two design choices for builtin support. Suppose we have files a.cpp and b.cpp, and each one includes header.h, generated by some action. Dependency graph created by classic jam would look like:

```

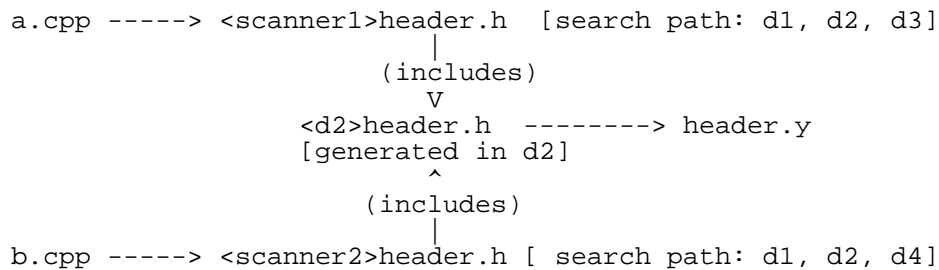
a.cpp -----> <scanner1>header.h [search path: d1, d2, d3]
                <d2>header.h -----> header.y
                [generated in d2]
b.cpp -----> <scanner2>header.h [ search path: d1, d2, d4]

```

In this case, Jam thinks all header.h target are not realated. The right dependency graph might be:



or



The first alternative was used for some time. The problem however is: what include paths should be used when scanning header.h? The second alternative was suggested by Matt Armstrong. It has similiar effect: add targets which depend on <scanner1>header.h will also depend on <d2>header.h. But now we have two different target with two different scanners, and those targets can be scanned independently. The problem of first alternative is avoided, so the second alternative is implemented now.

The second sub-requirements is that targets generated to "bin" directory are handled as well. Boost.Build implements semi-automatic approach. When compiling C++ files the process is:

1. The main target to which compiled file belongs is found.
2. All other main targets that the found one depends on are found. Those include main target which are used as sources, or present as values of "dependency" features.
3. All directories where files belonging to those main target will be generated are added to the include path.

After this is done, dependencies are found by the approach explained previously.

Note that if a target uses generated headers from other main target, that main target should be explicitly specified as dependency property. It would be better to lift this requirement, but it seems not very problematic in practice.

For target types other than C++, adding of include paths must be implemented anew.

Proper detection of dependencies from generated files

Suppose file "a.cpp" includes "a.h" and both are generated by some action. Note that classic jam has two stages. In first stage dependency graph graph is build and actions which should be run are determined. In second stage the actions are executed. Initially, neither file exists, so the include is not found. As the result, jam might attempt to compile a.cpp before creating a.h, and compilation will fail.

The solution in Boost.Jam is to perform additional dependency scans after targets are updated. This break separation

between build stages in jam — which some people consider a good thing — but I'm not aware of any better solution.

In order to understand the rest of this section, you better read some details about jam dependency scanning, available at this link.

Whenever a target is updated, Boost.Jam rescans it for includes. Consider this graph, created before any actions are run.

```
A -----> C ----> C.pro
      /
B --/      C-includes ----> D
```

Both A and B have dependency on C and C-includes (the latter dependency is not shown). Say during building we've tried to create A, then tried to create C and successfully created C.

In that case, the set of includes in C might well have changed. We do not bother to detect precisely which includes were added or removed. Instead we create another internal node C-includes-2. Then we determine what actions should be run to update the target. In fact this mean that we perform logic of first stage while already executing stage.

After actions for C-includes-2 are determined, we add C-includes-2 to the list of A's dependents, and stage 2 proceeds as usual. Unfortunately, we can't do the same with target B, since when it's not visited, C target does not know B depends on it. So, we add a flag to C which tells and it was rescanned. When visiting B target, the flag is notices and C-includes-2 will be added to the list of B's dependencies.

Note also that internal nodes are sometimes updated too. Consider this dependency graph:

```
a.o ----> a.cpp
           a.cpp-includes --> a.h (scanned)
                               a.h-includes -----> a.h (generated)
                                                                    |
a.pro <-----+-----
```

Here, out handling of generated headers come into play. Say that a.h exists but is out of date with respect to "a.pro", then "a.h (generated)" and "a.h-includes" will be marking for updating, but "a.h (scanned)" won't be marked. We have to rescan "a.h" file after it's created, but since "a.h (generated)" has no scanner associated with it, it's only possible to rescan "a.h" after "a.h-includes" target was updated.

The above consideration lead to decision that we'll rescan a target whenever it's updated, no matter if this target is internal or not.

Warning

The remainder of this document is not intended to be read at all. This will be rearranged in future.

File targets

As described above, file targets corresponds to files that Boost.Build manages. User's may be concerned about file targets in three ways: when declaring file target types, when declaring transformations between types, and when determining where file target will be placed. File targets can also be connected with actions, that determine how the target is created. Both file targets and actions are implemented in the `virtual-target` module.

Types

A file target can be given a file, which determines what transformations can be applied to the file. The

`type.register` rule declares new types. File type can also be assigned a scanner, which is used to find implicit dependencies. See "dependency scanning" [[link?](#)] below.

Target paths

To distinguish targets build with different properties, they are put in different directories. Rules for determining target paths are given below:

1. All targets are placed under directory corresponding to the project where they are defined.
2. Each non free, non incidental property cause an additional element to be added to the target path. That element has the form `<feature-name>-<feature-value>` for ordinary features and `<feature-value>` for implicit ones. [Note about composite features].
3. If the set of free, non incidental properties is different from the set of free, non incidental properties for the project in which the main target that uses the target is defined, a part of the form `main_target-<name>` is added to the target path. **Note:**It would be nice to completely track free features also, but this appears to be complex and not extremely needed.

For example, we might have these paths:

```
debug/optimization-off  
debug/main-target-a
```