

Richard Purdie, Linux Foundation
<**richard.purdie@linuxfoundation.org**>

by Richard Purdie
Copyright © 2010-2013 Linux Foundation

Permission is granted to copy, distribute and/or modify this document under the terms of the Creative Commons Attribution-Share Alike 2.0 UK: England & Wales [<http://creativecommons.org/licenses/by-sa/2.0/uk/>] as published by Creative Commons.

Manual Notes

- This version of The Yocto Project Reference Manual is for the 1.2.2 release of the Yocto Project. To be sure you have the latest version of the manual for this release, go to the Yocto Project documentation page [<http://www.yoctoproject.org/documentation>] and select the manual from that site. Manuals from the site are more up-to-date than manuals derived from the Yocto Project released TAR files.
- If you located this manual through a web search, the version of the manual might not be the one you want (e.g. the search might have returned a manual much older than the Yocto Project version with which you are working). You can see all Yocto Project major releases by visiting the Releases [<https://wiki.yoctoproject.org/wiki/Releases>] page. If you need a version of this manual for a different Yocto Project release, visit the Yocto Project documentation page [<http://www.yoctoproject.org/documentation>] and select the manual set by using the "ACTIVE RELEASES DOCUMENTATION" or "DOCUMENTS ARCHIVE" pull-down menus.
- To report any inaccuracies or problems with this manual, send an email to the Yocto Project discussion group at yocto@yoctoproject.com or log into the freenode #yocto channel.

Table of Contents

1. Introduction	1
1.1. Introduction	1
1.2. Documentation Overview	1
1.3. System Requirements	1
1.4. Obtaining the Yocto Project	1
1.5. Development Checkouts	2
2. Using the Yocto Project	3
2.1. Running a Build	3
2.1.1. Build Overview	3
2.1.2. Building an Image Using GPL Components	3
2.2. Installing and Using the Result	3
2.3. Debugging Build Failures	3
2.3.1. Task Failures	4
2.3.2. Running Specific Tasks	4
2.3.3. Dependency Graphs	4
2.3.4. General BitBake Problems	4
2.3.5. Building with No Dependencies	5
2.3.6. Variables	5
2.3.7. Recipe Logging Mechanisms	5
2.3.8. Other Tips	6
3. Technical Details	7
3.1. Yocto Project Components	7
3.1.1. BitBake	7
3.1.2. Metadata (Recipes)	8
3.1.3. Classes	8
3.1.4. Configuration	8
3.2. Shared State Cache	8
3.2.1. Overall Architecture	8
3.2.2. Checksums (Signatures)	9
3.2.3. Shared State	10
3.2.4. Tips and Tricks	11
3.3. x32	12
3.3.1. Support	12
3.3.2. Future Development and Limitations	13
3.3.3. Using x32 Right Now	13
3.4. Licenses	13
3.4.1. Tracking License Changes	14
3.4.2. Enabling Commercially Licensed Recipes	15
4. Board Support Packages (BSP) - Developer's Guide	17
4.1. BSP Layers	17
4.2. Example Filesystem Layout	17
4.2.1. License Files	19
4.2.2. README File	19
4.2.3. README.sources File	19
4.2.4. Pre-built User Binaries	19
4.2.5. Layer Configuration File	20
4.2.6. Hardware Configuration Options	20
4.2.7. Miscellaneous Recipe Files	21
4.2.8. Core Recipe Files	21
4.2.9. Display Support Files	21
4.2.10. Linux Kernel Configuration	22
4.3. Requirements and Recommendations for Released BSPs	24
4.3.1. Released BSP Requirements	24
4.3.2. Released BSP Recommendations	25
4.4. Customizing a Recipe for a BSP	26
4.5. BSP Licensing Considerations	26
4.6. Using the Yocto Project's BSP Tools	27
4.6.1. Common Features	27
4.6.2. Creating a new BSP Layer Using the yocto-bsp Script	29
4.6.3. Managing Kernel Patches and Config Items with yocto-kernel	31
5. Platform Development with the Yocto Project	34

5.1. Application Development Using the Yocto Project	34
5.1.1. External Development Using the Meta-Toolchain	34
5.1.2. External Development Using the Eclipse Plug-in	34
5.1.3. External Development Using the QEMU Emulator	35
5.1.4. Development Using Yocto Project Directly	35
5.1.5. Development Within a Development Shell	36
5.1.6. Development Within Yocto Project for a Package that Uses an External SCM	36
5.2. Debugging With the GNU Project Debugger (GDB) Remotely	37
5.2.1. Launching Gdbserver on the Target	37
5.2.2. Launching GDB on the Host Computer	38
5.3. Profiling with OProfile	39
5.3.1. Profiling on the Target	40
5.3.2. Using OProfileUI	40
A. Reference: Directory Structure	43
A.1. Top level core components	43
A.1.1. bitbake/	43
A.1.2. build/	43
A.1.3. documentation	43
A.1.4. meta/	43
A.1.5. meta-demoapps/	43
A.1.6. meta-rt/	43
A.1.7. meta-skeleton/	43
A.1.8. scripts/	43
A.1.9. oe-init-build-env	44
A.1.10. LICENSE, README, and README.hardware	44
A.2. The Build Directory - build/	44
A.2.1. build/pseudodone	44
A.2.2. build/conf/local.conf	44
A.2.3. build/conf/bblayers.conf	44
A.2.4. build/conf/sanity_info	44
A.2.5. build/downloads/	44
A.2.6. build/sstate-cache/	44
A.2.7. build/tmp/	45
A.2.8. build/tmp/buildstats/	45
A.2.9. build/tmp/cache/	45
A.2.10. build/tmp/deploy/	45
A.2.11. build/tmp/deploy/deb/	45
A.2.12. build/tmp/deploy/rpm/	45
A.2.13. build/tmp/deploy/images/	45
A.2.14. build/tmp/deploy/ipk/	45
A.2.15. build/tmp/sysroots/	45
A.2.16. build/tmp/stamps/	45
A.2.17. build/tmp/log/	45
A.2.18. build/tmp/pkgdata/	46
A.2.19. build/tmp/work/	46
A.3. The Metadata - meta/	46
A.3.1. meta/classes/	46
A.3.2. meta/conf/	46
A.3.3. meta/conf/machine/	46
A.3.4. meta/conf/distro/	46
A.3.5. meta/recipes-bsp/	47
A.3.6. meta/recipes-connectivity/	47
A.3.7. meta/recipes-core/	47
A.3.8. meta/recipes-devtools/	47
A.3.9. meta/recipes-extended/	47
A.3.10. meta/recipes-gnome/	47
A.3.11. meta/recipes-graphics/	47
A.3.12. meta/recipes-kernel/	47
A.3.13. meta/recipes-multimedia/	47
A.3.14. meta/recipes-qt/	47
A.3.15. meta/recipes-sato/	47
A.3.16. meta/recipes-support/	47
A.3.17. meta/site/	48
A.3.18. meta/recipes.txt/	48

B. Reference: BitBake	49
B.1. Parsing	49
B.2. Preferences and Providers	49
B.3. Dependencies	50
B.4. The Task List	50
B.5. Running a Task	50
B.6. BitBake Command Line	51
B.7. Fetchers	52
C. Reference: Classes	53
C.1. The base class - base.bbclass	53
C.2. Autotooled Packages - autotools.bbclass	53
C.3. Alternatives - update-alternatives.bbclass	53
C.4. Initscripts - update-rc.d.bbclass	54
C.5. Binary config scripts - binconfig.bbclass	54
C.6. Debian renaming - debian.bbclass	54
C.7. Pkg-config - pkgconfig.bbclass	54
C.8. Distribution of sources - src_distribute_local.bbclass	54
C.9. Perl modules - cpan.bbclass	54
C.10. Python extensions - distutils.bbclass	55
C.11. Developer Shell - devshell.bbclass	55
C.12. Packaging - package*.bbclass	55
C.13. Building kernels - kernel.bbclass	55
C.14. Creating images - image.bbclass and rootfs*.bbclass	56
C.15. Host System sanity checks - sanity.bbclass	56
C.16. Generated output quality assurance checks - insane.bbclass	56
C.17. Autotools configuration data cache - siteinfo.bbclass	57
C.18. Adding Users - useradd.bbclass	57
C.19. Using External Source - externalsrc.bbclass	57
C.20. Other Classes	58
D. Reference: Images	59
E. Reference: Features	61
E.1. Distro	61
E.2. Machine	61
E.3. Reference: Images	62
F. Reference: Variables Glossary	63
G. Reference: Variable Context	81
G.1. Configuration	81
G.1.1. Distribution (Distro)	81
G.1.2. Machine	81
G.1.3. Local	81
G.2. Recipes	82
G.2.1. Required	82
G.2.2. Dependencies	82
G.2.3. Paths	82
G.2.4. Extra Build Information	82
H. FAQ	84
I. Contributing to the Yocto Project	89
I.1. Introduction	89
I.2. Tracking Bugs	89
I.3. Mailing lists	89
I.4. Internet Relay Chat (IRC)	89
I.5. Links	89
I.6. Contributions	90

Chapter 1. Introduction

1.1. Introduction

This manual provides reference information for the current release of the Yocto Project. The Yocto Project is an open-source collaboration project focused on embedded Linux developers. Amongst other things, the Yocto Project uses the Poky build tool to construct complete Linux images. You can find complete introductory and getting started information on the Yocto Project by reading the Yocto Project Quick Start [<http://www.yoctoproject.org/docs/1.2.2/yocto-project-qs/yocto-project-qs.html>]. For task-based information using the Yocto Project, see The Yocto Project Development Manual [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html>]. You can also find lots of information on the Yocto Project on the Yocto Project website [<http://www.yoctoproject.org>].

1.2. Documentation Overview

This reference manual consists of the following:

- Using the Yocto Project: This chapter provides an overview of the components that make up the Yocto Project followed by information about debugging images created in the Yocto Project.
- Technical Details: This chapter describes fundamental Yocto Project components as well as an explanation behind how the Yocto Project uses shared state (sstate) cache to speed build time.
- Board Support Packages (BSP) - Developer's Guide: This chapter describes the example filesystem layout for BSP development and the click-through licensing scheme.
- Platform Development With the Yocto Project: This chapter describes application development, debugging, and profiling using the Yocto Project.
- Reference: Directory Structure: This appendix describes the directory structure of the Yocto Project files. The Yocto Project files represent the file structure or Git repository created as a result of setting up the Yocto Project on your host development system.
- Reference: BitBake: This appendix provides an overview of the BitBake tool and its role within the Yocto Project.
- Reference: Classes: This appendix describes the classes used in the Yocto Project.
- Reference: Images: This appendix describes the standard images that the Yocto Project supports.
- Reference: Features: This appendix describes mechanisms for creating distribution, machine, and image features during the build process using the Yocto Project.
- Reference: Variables Glossary: This appendix presents most Yocto Project variables. Entries describe the function of the variable and how to apply them.
- Reference: Variable Context: This appendix provides variable locality or context.
- Reference: FAQ: This appendix provides answers for commonly asked questions in the Yocto Project development environment.
- Reference: Contributing to the Yocto Project: This appendix provides guidance on how you can contribute back to the Yocto Project.

1.3. System Requirements

For Yocto Project system requirements, see the What You Need and How You Get It [<http://www.yoctoproject.org/docs/1.2.2/yocto-project-qs/yocto-project-qs.html#resources>] section in the Yocto Project Quick Start.

1.4. Obtaining the Yocto Project

The Yocto Project development team makes the Yocto Project available through a number of methods:

- Releases: Stable, tested releases are available through <http://downloads.yoctoproject.org/releases/yocto/>.
- Nightly Builds: These releases are available at <http://autobuilder.yoctoproject.org/nightly>. These builds include Yocto Project releases, meta-toolchain tarballs, and experimental builds.
- Yocto Project Website: You can find releases of the Yocto Project and supported BSPs at the Yocto Project website [<http://www.yoctoproject.org>]. Along with these downloads, you can find lots of other information at this site.

1.5. Development Checkouts

Development using the Yocto Project requires a local copy of the Yocto Project files. You can get these files by downloading a Yocto Project release tarball and unpacking it, or by establishing a Git repository of the files. For information on both these methods, see the "Getting Setup [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#getting-setup>]" section in The Yocto Project Development Manual.

Chapter 2. Using the Yocto Project

This chapter describes common usage for the Yocto Project. The information is introductory in nature as other manuals in the Yocto Project provide more details on how to use the Yocto Project.

2.1. Running a Build

You can find general information on how to build an image using the Yocto Project in the "Building an Image [<http://www.yoctoproject.org/docs/1.2.2/yocto-project-qs/yocto-project-qs.html#building-image>]" section of The Yocto Project Quick Start. This section provides a summary of the build process and provides information for less obvious aspects of the build process.

2.1.1. Build Overview

The first thing you need to do is set up the Yocto Project build environment by sourcing the environment setup script as follows:

```
$ source oe-init-build-env [build_dir]
```

The `build_dir` is optional and specifies the directory Yocto Project uses for the build. If you do not specify a build directory it defaults to build in your current working directory. A common practice is to use a different build directory for different targets. For example, `~/build/x86` for a `qemux86` target, and `~/build/arm` for a `qemuarm` target. See `oe-init-build-env` for more information on this script.

Once the Yocto Project build environment is set up, you can build a target using:

```
$ bitbake <target>
```

The target is the name of the recipe you want to build. Common targets are the images in `meta/recipes-core/images`, `/meta/recipes-sato/images`, etc. all found in the Yocto Project files. Or, the target can be the name of a recipe for a specific piece of software such as `busybox`. For more details about the images the Yocto Project supports, see the "Reference: Images" appendix.

Note

Building an image without GNU Public License Version 3 (GPLv3) components is only supported for minimal and base images. See the "Reference: Images" appendix for more information.

2.1.2. Building an Image Using GPL Components

When building an image using GPL components, you need to maintain your original settings and not switch back and forth applying different versions of the GNU Public License. If you rebuild using different versions of GPL, dependency errors might occur due to some components not being rebuilt.

2.2. Installing and Using the Result

Once an image has been built, it often needs to be installed. The images and kernels built by the Yocto Project are placed in the build directory in `tmp/deploy/images`. For information on how to run pre-built images such as `qemux86` and `qemuarm`, see the "Using Pre-Built Binaries and QEMU [<http://www.yoctoproject.org/docs/1.2.2/yocto-project-qs/yocto-project-qs.html#using-pre-built>]" section in the Yocto Project Quick Start. For information about how to install these images, see the documentation for your particular board/machine.

2.3. Debugging Build Failures

The exact method for debugging Yocto Project build failures depends on the nature of the problem and on the system's area from which the bug originates. Standard debugging practices such as comparison against the last known working version with examination of the changes and the re-

application of steps to identify the one causing the problem are valid for Yocto Project just as they are for any other system. Even though it is impossible to detail every possible potential failure, this section provides some general tips to aid in debugging.

2.3.1. Task Failures

The log file for shell tasks is available in `${WORKDIR}/temp/log.do_taskname.pid`. For example, the compile task for the QEMU minimal image for the x86 machine (qemux86) might be `tmp/work/qemux86-poky-linux/core-image-minimal-1.0-r0/temp/log.do_compile.20830`. To see what BitBake runs to generate that log, look at the corresponding `run.do_taskname.pid` file located in the same directory.

Presently, the output from Python tasks is sent directly to the console.

2.3.2. Running Specific Tasks

Any given package consists of a set of tasks. The standard BitBake behavior in most cases is: fetch, unpack, patch, configure, compile, install, package, package_write, and build. The default task is build and any tasks on which it depends build first. Some tasks exist, such as devshell, that are not part of the default build chain. If you wish to run a task that is not part of the default build chain, you can use the `-c` option in BitBake as follows:

```
$ bitbake matchbox-desktop -c devshell
```

If you wish to rerun a task, use the `-f` force option. For example, the following sequence forces recompilation after changing files in the working directory.

```
$ bitbake matchbox-desktop
.
.
[make some changes to the source code in the working directory]
.
.
$ bitbake matchbox-desktop -c compile -f
$ bitbake matchbox-desktop
```

This sequence first builds `matchbox-desktop` and then recompiles it. The last command reruns all tasks (basically the packaging tasks) after the compile. BitBake recognizes that the compile task was rerun and therefore understands that the other tasks also need to be run again.

You can view a list of tasks in a given package by running the `listtasks` task as follows:

```
$ bitbake matchbox-desktop -c listtasks
```

The results are in the file `${WORKDIR}/temp/log.do_listtasks`.

2.3.3. Dependency Graphs

Sometimes it can be hard to see why BitBake wants to build some other packages before a given package you have specified. The `bitbake -g targetname` command creates the `depends.dot` and `task-depends.dot` files in the current directory. These files show the package and task dependencies and are useful for debugging problems. You can use the `bitbake -g -u depexp targetname` command to display the results in a more human-readable form.

2.3.4. General BitBake Problems

You can see debug output from BitBake by using the `-D` option. The debug output gives more information about what BitBake is doing and the reason behind it. Each `-D` option you use increases the logging level. The most common usage is `-DDD`.

The output from `bitbake -DDD -v targetname` can reveal why BitBake chose a certain version of a package or why BitBake picked a certain provider. This command could also help you in a situation where you think BitBake did something unexpected.

2.3.5. Building with No Dependencies

If you really want to build a specific `.bb` file, you can use the command form `bitbake -b <somepath/somefile.bb>`. This command form does not check for dependencies so you should use it only when you know its dependencies already exist. You can also specify fragments of the filename. In this case, BitBake checks for a unique match.

2.3.6. Variables

The `-e` option dumps the resulting environment for either the configuration (no package specified) or for a specific package when specified; or `-b recipename` to show the environment from parsing a single recipe file only.

2.3.7. Recipe Logging Mechanisms

Best practices exist while writing recipes that both log build progress and act on build conditions such as warnings and errors. Both Python and Bash language bindings exist for the logging mechanism:

- Python: For Python functions, BitBake supports several loglevels: `bb.fatal`, `bb.error`, `bb.warn`, `bb.note`, `bb.plain`, and `bb.debug`.
- Bash: For Bash functions, the same set of loglevels exist and are accessed with a similar syntax: `bbfatal`, `bberror`, `bbwarn`, `bbnote`, `bbplain`, and `bbdebug`.

For guidance on how logging is handled in both Python and Bash recipes, see the `logging.bbclass` file in the `meta/classes` directory of the Yocto Project files.

2.3.7.1. Logging With Python

When creating recipes using Python and inserting code that handles build logs keep in mind the goal is to have informative logs while keeping the console as "silent" as possible. Also, if you want status messages in the log use the "debug" loglevel.

Following is an example written in Python. The code handles logging for a function that determines the number of tasks needed to be run:

```
python do_listtasks() {
    bb.debug(2, "Starting to figure out the task list")
    if noteworthy_condition:
        bb.note("There are 47 tasks to run")
    bb.debug(2, "Got to point xyz")
    if warning_trigger:
        bb.warn("Detected warning_trigger, this might be a problem later.")
    if recoverable_error:
        bb.error("Hit recoverable_error, you really need to fix this!")
    if fatal_error:
        bb.fatal("fatal_error detected, unable to print the task list")
    bb.plain("The tasks present are abc")
    bb.debug(2, "Finished figuring out the tasklist")
}
```

2.3.7.2. Logging With Bash

When creating recipes using Bash and inserting code that handles build logs you have the same goals - informative with minimal console output. The syntax you use for recipes written in Bash is similar to that of recipes written in Python described in the previous section.

Following is an example written in Bash. The code logs the progress of the `do_my_function` function.

```
do_my_function() {
    bbdebug 2 "Running do_my_function"
    if [ exceptional_condition ]; then
        bbnote "Hit exceptional_condition"
    fi
    bbdebug 2 "Got to point xyz"
    if [ warning_trigger ]; then
        bbwarn "Detected warning_trigger, this might cause a problem later."
    fi
    if [ recoverable_error ]; then
        bberror "Hit recoverable_error, correcting"
    fi
    if [ fatal_error ]; then
        bbfatal "fatal_error detected"
    fi
    bbdebug 2 "Completed do_my_function"
}
```

2.3.8. Other Tips

Here are some other tips that you might find useful:

- When adding new packages, it is worth watching for undesirable items making their way into compiler command lines. For example, you do not want references to local system files like `/usr/lib/` or `/usr/include/`.
- If you want to remove the psplash boot splashscreen, add `psplash=false` to the kernel command line. Doing so prevents psplash from loading and thus allows you to see the console. It is also possible to switch out of the splashscreen by switching the virtual console (e.g. `Fn+Left` or `Fn+Right` on a Zaurus).

Chapter 3. Technical Details

This chapter provides technical details for various parts of the Yocto Project. Currently, topics include Yocto Project components and shared state (sstate) cache.

3.1. Yocto Project Components

The BitBake task executor together with various types of configuration files form the Yocto Project core. This section overviews the BitBake task executor and the configuration files by describing what they are used for and how they interact.

BitBake handles the parsing and execution of the data files. The data itself is of various types:

- Recipes: Provides details about particular pieces of software
- Class Data: An abstraction of common build information (e.g. how to build a Linux kernel).
- Configuration Data: Defines machine-specific settings, policy decisions, etc. Configuration data acts as the glue to bind everything together.

For more information on data, see the "Yocto Project Terms [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#yocto-project-terms>]" section in the Yocto Project Development Manual.

BitBake knows how to combine multiple data sources together and refers to each data source as a layer. For information on layers, see the "Understanding and Creating Layers [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#understanding-and-creating-layers>]" section of the Yocto Project Development Manual.

Following are some brief details on these core components. For more detailed information on these components see the "Reference: Directory Structure" appendix.

3.1.1. BitBake

BitBake is the tool at the heart of the Yocto Project and is responsible for parsing the metadata, generating a list of tasks from it, and then executing those tasks. To see a list of the options BitBake supports, use the following help command:

```
$ bitbake --help
```

The most common usage for BitBake is `bitbake <packagename>`, where `packagename` is the name of the package you want to build (referred to as the "target" in this manual). The target often equates to the first part of a `.bb` filename. So, to run the `matchbox-desktop_1.2.3.bb` file, you might type the following:

```
$ bitbake matchbox-desktop
```

Several different versions of `matchbox-desktop` might exist. BitBake chooses the one selected by the distribution configuration. You can get more details about how BitBake chooses between different target versions and providers in the "Preferences and Providers" section.

BitBake also tries to execute any dependent tasks first. So for example, before building `matchbox-desktop`, BitBake would build a cross compiler and `eglibc` if they had not already been built.

Note

This release of the Yocto Project does not support the `glibc` GNU version of the Unix standard C library. By default, the Yocto Project builds with `eglibc`.

A useful BitBake option to consider is the `-k` or `--continue` option. This option instructs BitBake to try and continue processing the job as much as possible even after encountering an error. When an error occurs, the target that failed and those that depend on it cannot be remade. However, when you use this option other dependencies can still be processed.

3.1.2. Metadata (Recipes)

The `.bb` files are usually referred to as "recipes." In general, a recipe contains information about a single piece of software. The information includes the location from which to download the source patches (if any are needed), which special configuration options to apply, how to compile the source files, and how to package the compiled output.

The term "package" can also be used to describe recipes. However, since the same word is used for the packaged output from the Yocto Project (i.e. `.ipk` or `.deb` files), this document avoids using the term "package" when referring to recipes.

3.1.3. Classes

Class files (`.bbclass`) contain information that is useful to share between metadata files. An example is the Autotools class, which contains common settings for any application that Autotools uses. The "Reference: Classes" appendix provides details about common classes and how to use them.

3.1.4. Configuration

The configuration files (`.conf`) define various configuration variables that govern the Yocto Project build process. These files fall into several areas that define machine configuration options, distribution configuration options, compiler tuning options, general common configuration options and user configuration options (`local.conf`, which is found in the Yocto Project files build directory).

3.2. Shared State Cache

By design, the Yocto Project build system builds everything from scratch unless BitBake can determine that parts don't need to be rebuilt. Fundamentally, building from scratch is attractive as it means all parts are built fresh and there is no possibility of stale data causing problems. When developers hit problems, they typically default back to building from scratch so they know the state of things from the start.

Building an image from scratch is both an advantage and a disadvantage to the process. As mentioned in the previous paragraph, building from scratch ensures that everything is current and starts from a known state. However, building from scratch also takes much longer as it generally means rebuilding things that don't necessarily need rebuilt.

The Yocto Project implements shared state code that supports incremental builds. The implementation of the shared state code answers the following questions that were fundamental roadblocks within the Yocto Project incremental build support system:

- What pieces of the system have changed and what pieces have not changed?
- How are changed pieces of software removed and replaced?
- How are pre-built components that don't need to be rebuilt from scratch used when they are available?

For the first question, the build system detects changes in the "inputs" to a given task by creating a checksum (or signature) of the task's inputs. If the checksum changes, the system assumes the inputs have changed and the task needs to be rerun. For the second question, the shared state (`sstate`) code tracks which tasks add which output to the build process. This means the output from a given task can be removed, upgraded or otherwise manipulated. The third question is partly addressed by the solution for the second question assuming the build system can fetch the `sstate` objects from remote locations and install them if they are deemed to be valid.

The rest of this section goes into detail about the overall incremental build architecture, the checksums (signatures), shared state, and some tips and tricks.

3.2.1. Overall Architecture

When determining what parts of the system need to be built, BitBake uses a per-task basis and does not use a per-recipe basis. You might wonder why using a per-task basis is preferred over a per-recipe

basis. To help explain, consider having the IPK packaging backend enabled and then switching to DEB. In this case, `do_install` and `do_package` output are still valid. However, with a per-recipe approach, the build would not include the `.deb` files. Consequently, you would have to invalidate the whole build and rerun it. Rerunning everything is not the best situation. Also in this case, the core must be "taught" much about specific tasks. This methodology does not scale well and does not allow users to easily add new tasks in layers or as external recipes without touching the packaged-staging core.

3.2.2. Checksums (Signatures)

The shared state code uses a checksum, which is a unique signature of a task's inputs, to determine if a task needs to be run again. Because it is a change in a task's inputs that triggers a rerun, the process needs to detect all the inputs to a given task. For shell tasks, this turns out to be fairly easy because the build process generates a "run" shell script for each task and it is possible to create a checksum that gives you a good idea of when the task's data changes.

To complicate the problem, there are things that should not be included in the checksum. First, there is the actual specific build path of a given task - the `WORKDIR`. It does not matter if the working directory changes because it should not affect the output for target packages. Also, the build process has the objective of making native/cross packages relocatable. The checksum therefore needs to exclude `WORKDIR`. The simplistic approach for excluding the working directory is to set `WORKDIR` to some fixed value and create the checksum for the "run" script.

Another problem results from the "run" scripts containing functions that might or might not get called. The incremental build solution contains code that figures out dependencies between shell functions. This code is used to prune the "run" scripts down to the minimum set, thereby alleviating this problem and making the "run" scripts much more readable as a bonus.

So far we have solutions for shell scripts. What about python tasks? The same approach applies even though these tasks are more difficult. The process needs to figure out what variables a python function accesses and what functions it calls. Again, the incremental build solution contains code that first figures out the variable and function dependencies, and then creates a checksum for the data used as the input to the task.

Like the `WORKDIR` case, situations exist where dependencies should be ignored. For these cases, you can instruct the build process to ignore a dependency by using a line like the following:

```
PACKAGE_ARCHS[vardepsexclude] = "MACHINE"
```

This example ensures that the `PACKAGE_ARCHS` variable does not depend on the value of `MACHINE`, even if it does reference it.

Equally, there are cases where we need to add dependencies BitBake is not able to find. You can accomplish this by using a line like the following:

```
PACKAGE_ARCHS[vardeps] = "MACHINE"
```

This example explicitly adds the `MACHINE` variable as a dependency for `PACKAGE_ARCHS`.

Consider a case with inline python, for example, where BitBake is not able to figure out dependencies. When running in debug mode (i.e. using `-DDD`), BitBake produces output when it discovers something for which it cannot figure out dependencies. The Yocto Project team has currently not managed to cover those dependencies in detail and is aware of the need to fix this situation.

Thus far, this section has limited discussion to the direct inputs into a task. Information based on direct inputs is referred to as the "basehash" in the code. However, there is still the question of a task's indirect inputs - the things that were already built and present in the build directory. The checksum (or signature) for a particular task needs to add the hashes of all the tasks on which the particular task depends. Choosing which dependencies to add is a policy decision. However, the effect is to generate a master checksum that combines the basehash and the hashes of the task's dependencies.

At the code level, there are a variety of ways both the basehash and the dependent task hashes can be influenced. Within the BitBake configuration file, we can give BitBake some extra information to

help it construct the basehash. The following statements effectively result in a list of global variable dependency excludes - variables never included in any checksum:

```
BB_HASHBASE_WHITELIST ?= "TMPDIR FILE_PATH PWD BB_TASKHASH BBPATH"
BB_HASHBASE_WHITELIST += "DL_DIR SSTATE_DIR THISDIR FILESEXTRAPATHS"
BB_HASHBASE_WHITELIST += "FILE_DIRNAME HOME LOGNAME SHELL TERM USER"
BB_HASHBASE_WHITELIST += "FILESPATH USERNAME STAGING_DIR_HOST STAGING_DIR_TARGET"
```

The previous example actually excludes WORKDIR since it is actually constructed as a path within TMPDIR, which is on the whitelist.

The rules for deciding which hashes of dependent tasks to include through dependency chains are more complex and are generally accomplished with a python function. The code in meta/lib/oe/sstatesig.py shows two examples of this and also illustrates how you can insert your own policy into the system if so desired. This file defines the two basic signature generators OE-Core uses: "OEBasic" and "OEBasicHash". By default, there is a dummy "noop" signature handler enabled in BitBake. This means that behavior is unchanged from previous versions. OE-Core uses the "OEBasic" signature handler by default through this setting in the bitbake.conf file:

```
BB_SIGNATURE_HANDLER ?= "OEBasic"
```

The "OEBasicHash" BB_SIGNATURE_HANDLER is the same as the "OEBasic" version but adds the task hash to the stamp files. This results in any metadata change that changes the task hash, automatically causing the task to be run again. This removes the need to bump PR values and changes to metadata automatically ripple across the build. Currently, this behavior is not the default behavior for OE-Core but is the default in poky.

It is also worth noting that the end result of these signature generators is to make some dependency and hash information available to the build. This information includes:

```
BB_BASEHASH_task-<taskname> - the base hashes for each task in the recipe
BB_BASEHASH_<filename:taskname> - the base hashes for each dependent task
BBHASHDEPS_<filename:taskname> - The task dependencies for each task
BB_TASKHASH - the hash of the currently running task
```

3.2.3. Shared State

Checksums and dependencies, as discussed in the previous section, solve half the problem. The other part of the problem is being able to use checksum information during the build and being able to reuse or rebuild specific components.

The shared state class (sstate.bbclass) is a relatively generic implementation of how to "capture" a snapshot of a given task. The idea is that the build process does not care about the source of a task's output. Output could be freshly built or it could be downloaded and unpacked from somewhere - the build process doesn't need to worry about its source.

There are two types of output, one is just about creating a directory in WORKDIR. A good example is the output of either do_install or do_package. The other type of output occurs when a set of data is merged into a shared directory tree such as the sysroot.

The Yocto Project team has tried to keep the details of the implementation hidden in sstate.bbclass. From a user's perspective, adding shared state wrapping to a task is as simple as this do_deploy example taken from do_deploy.bbclass:

```
DEPLOYDIR = "${WORKDIR}/deploy-${PN}"
SSTATETASKS += "do_deploy"
do_deploy[sstate-name] = "deploy"
do_deploy[sstate-inputdirs] = "${DEPLOYDIR}"
do_deploy[sstate-outputdirs] = "${DEPLOY_DIR_IMAGE}"
```

```
python do_deploy_setscene () {
    sstate_setscene(d)
}
addtask do_deploy_setscene
```

In the example, we add some extra flags to the task, a name field ("deploy"), an input directory where the task sends data, and the output directory where the data from the task should eventually be copied. We also add a `_setscene` variant of the task and add the task name to the `SSTATETASKS` list.

If you have a directory whose contents you need to preserve, you can do this with a line like the following:

```
do_package[sstate-plaindirs] = "${PKGDEST} ${PKGDEST}"
```

This method, as well as the following example, also works for multiple directories.

```
do_package[sstate-inputdirs] = "${PKGDESTWORK} ${SHLIBSWORKDIR}"
do_package[sstate-outputdirs] = "${PKGDATA_DIR} ${SHLIBSDIR}"
do_package[sstate-lockfile] = "${PACKAGELOCK}"
```

These methods also include the ability to take a lockfile when manipulating shared state directory structures since some cases are sensitive to file additions or removals.

Behind the scenes, the shared state code works by looking in `SSTATE_DIR` and `SSTATE_MIRRORS` for shared state files. Here is an example:

```
SSTATE_MIRRORS ?= "\
file:///.* http://someserver.tld/share/sstate/ \n \
file:///.* file:///some/local/dir/sstate/"
```

The shared state package validity can be detected just by looking at the filename since the filename contains the task checksum (or signature) as described earlier in this section. If a valid shared state package is found, the build process downloads it and uses it to accelerate the task.

The build processes uses the `*_setscene` tasks for the task acceleration phase. BitBake goes through this phase before the main execution code and tries to accelerate any tasks for which it can find shared state packages. If a shared state package for a task is available, the shared state package is used. This means the task and any tasks on which it is dependent are not executed.

As a real world example, the aim is when building an IPK-based image, only the `do_package_write_ipk` tasks would have their shared state packages fetched and extracted. Since the `sysroot` is not used, it would never get extracted. This is another reason why a task-based approach is preferred over a recipe-based approach, which would have to install the output from every task.

3.2.4. Tips and Tricks

The code in the Yocto Project that supports incremental builds is not simple code. This section presents some tips and tricks that help you work around issues related to shared state code.

3.2.4.1. Debugging

When things go wrong, debugging needs to be straightforward. Because of this, the Yocto Project team included strong debugging tools:

- Whenever a shared state package is written out, so is a corresponding `.sinfo` file. This practice results in a pickled python database of all the metadata that went into creating the hash for a given shared state package.

- If BitBake is run with the `--dump-signatures` (or `-S`) option, BitBake dumps out `.siginfo` files in the stamp directory for every task it would have executed instead of building the specified target package.
- There is a `bitbake-diffsigs` command that can process these `.siginfo` files. If one file is specified, it will dump out the dependency information in the file. If two files are specified, it will compare the two files and dump out the differences between the two. This allows the question of "What changed between X and Y?" to be answered easily.

3.2.4.2. Invalidating Shared State

The shared state code uses checksums and shared state memory cache to avoid unnecessarily rebuilding tasks. As with all schemes, this one has some drawbacks. It is possible that you could make implicit changes that are not factored into the checksum calculation, but do affect a task's output. A good example is perhaps when a tool changes its output. Let's say that the output of `rpmdeps` needed to change. The result of the change should be that all the "package", "package_write_rpm", and "package_deploy-rpm" shared state cache items would become invalid. But, because this is a change that is external to the code and therefore implicit, the associated shared state cache items do not become invalidated. In this case, the build process would use the cached items rather than running the task again. Obviously, these types of implicit changes can cause problems.

To avoid these problems during the build, you need to understand the effects of any change you make. Note that any changes you make directly to a function automatically are factored into the checksum calculation and thus, will invalidate the associated area of sstate cache. You need to be aware of any implicit changes that are not obvious changes to the code and could affect the output of a given task. Once you are aware of such a change, you can take steps to invalidate the cache and force the task to run. The step to take is as simple as changing a function's comments in the source code. For example, to invalidate package shared state files, change the comment statements of `do_package` or the comments of one of the functions it calls. The change is purely cosmetic, but it causes the checksum to be recalculated and forces the task to be run again.

Note

For an example of a commit that makes a cosmetic change to invalidate a shared state, see this commit [<http://git.yoctoproject.org/cgit.cgi/poky/commit/meta/classes/package.bbclass?id=737f8bbb4f27b4837047cb9b4fbfe01dfde36d54>].

3.3. x32

x32 is a new processor-specific Application Binary Interface (psABI) for x86_64. An ABI defines the calling conventions between functions in a processing environment. The interface determines what registers are used and what the sizes are for various C data types.

Some processing environments prefer using 32-bit applications even when running on Intel 64-bit platforms. Consider the i386 psABI, which is a very old 32-bit ABI for Intel 64-bit platforms. The i386 psABI does not provide efficient use and access of the Intel 64-bit processor resources, leaving the system underutilized. Now consider the x86_64 psABI. This ABI is newer and uses 64-bits for data sizes and program pointers. The extra bits increase the footprint size of the programs, libraries, and also increases the memory and file system size requirements. Executing under the x32 psABI enables user programs to utilize CPU and system resources more efficiently while keeping the memory footprint of the applications low. Extra bits are used for registers but not for addressing mechanisms.

3.3.1. Support

While the x32 psABI specifications are not fully finalized, this Yocto Project release supports current development specifications of x32 psABI. As of this release of the Yocto Project, x32 psABI support exists as follows:

- You can create packages and images in x32 psABI format on x86_64 architecture targets.
- You can use the x32 psABI support through the meta-x32 layer on top of the OE-core/Yocto layer.
- The toolchain from the experimental/meta-x32 layer is used for building x32 psABI program binaries.
- You can successfully build many recipes with the x32 toolchain.

- You can create and boot `core-image-minimal` and `core-image-sato` images.

3.3.2. Future Development and Limitations

As of this Yocto Project release, the x32 psABI kernel and library interfaces specifications are not finalized.

Future Plans for the x32 psABI in the Yocto Project include the following:

- Enhance and fix the few remaining recipes so they work with and support x32 toolchains.
- Enhance RPM Package Manager (RPM) support for x32 binaries.
- Support larger images.
- Integrate x32 recipes, toolchain, and kernel changes from `experimental/meta-x32` into `OE-core`.

3.3.3. Using x32 Right Now

Despite the fact the x32 psABI support is in development state for this release of the Yocto Project, you can follow these steps to use the x32 psABI:

- Add the `experimental/meta-x32` layer to your local Yocto Project Build Directory [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#yocto-project-build-directory>]. You can find the `experimental/meta-x32` source repository at <http://git.yoctoproject.org>.
- Edit your `conf/bblayers.conf` file so that it includes the `meta-x32`. Here is an example:

```
BBLAYERS ?= " \
    /home/nitin/prj/poky.git/meta \
    /home/nitin/prj/poky.git/meta-yocto \
    /home/nitin/prj/meta-x32.git \
"
```

- Enable the x32 psABI tuning file for x86_64 machines by editing the `conf/local.conf` like this:

```
MACHINE = "qemux86-64"
DEFAULTTUNE = "x86-64-x32"
baselib = "${@d.getVar('BASE_LIB_tune-' + (d.getVar('DEFAULTTUNE', True) \
    or 'INVALID'), True) or 'lib'}"
#MACHINE = "atom-pc"
#DEFAULTTUNE = "core2-64-x32"
```

- As usual, use BitBake to build an image that supports the x32 psABI. Here is an example:

```
$ bitake core-image-sato
```

- As usual, run your image using QEMU:

```
$ runqemu qemux86-64 core-image-sato
```

3.4. Licenses

This section describes the mechanism by which the Yocto Project build system tracks changes to licensing text. The section also describes how to enable commercially licensed recipes, which by default are disabled.

3.4.1. Tracking License Changes

The license of an upstream project might change in the future. In order to prevent these changes going unnoticed, the Yocto Project provides a `LIC_FILES_CHKSUM` variable to track changes to the license text. The checksums are validated at the end of the configure step, and if the checksums do not match, the build will fail.

3.4.1.1. Specifying the `LIC_FILES_CHKSUM` Variable

The `LIC_FILES_CHKSUM` variable contains checksums of the license text in the source code for the recipe. Following is an example of how to specify `LIC_FILES_CHKSUM`:

```
LIC_FILES_CHKSUM = "file://COPYING;md5=xxxx \
                    file://licfile1.txt;beginline=5;endline=29;md5=yyyy \
                    file://licfile2.txt;endline=50;md5=zzzz \
                    ..."
```

The Yocto Project uses the `S` variable as the default directory used when searching files listed in `LIC_FILES_CHKSUM`. The previous example employs the default directory.

You can also use relative paths as shown in the following example:

```
LIC_FILES_CHKSUM = "file://src/ls.c;beginline=5;endline=16;\
                    md5=bb14ed3c4cda583abc85401304b5cd4e"
LIC_FILES_CHKSUM = "file://../license.html;md5=5c94767cedb5d6987c902ac850ded2c6"
```

In this example, the first line locates a file in `${S}/src/ls.c`. The second line refers to a file in `WORKDIR`, which is the parent of `S`.

Note that this variable is mandatory for all recipes, unless the `LICENSE` variable is set to "CLOSED".

3.4.1.2. Explanation of Syntax

As mentioned in the previous section, the `LIC_FILES_CHKSUM` variable lists all the important files that contain the license text for the source code. It is possible to specify a checksum for an entire file, or a specific section of a file (specified by beginning and ending line numbers with the "beginline" and "endline" parameters, respectively). The latter is useful for source files with a license notice header, README documents, and so forth. If you do not use the "beginline" parameter, then it is assumed that the text begins on the first line of the file. Similarly, if you do not use the "endline" parameter, it is assumed that the license text ends with the last line of the file.

The "md5" parameter stores the md5 checksum of the license text. If the license text changes in any way as compared to this parameter then a mismatch occurs. This mismatch triggers a build failure and notifies the developer. Notification allows the developer to review and address the license text changes. Also note that if a mismatch occurs during the build, the correct md5 checksum is placed in the build log and can be easily copied to the recipe.

There is no limit to how many files you can specify using the `LIC_FILES_CHKSUM` variable. Generally, however, every project requires a few specifications for license tracking. Many projects have a "COPYING" file that stores the license information for all the source code files. This practice allows you to just track the "COPYING" file as long as it is kept up to date.

Tip

If you specify an empty or invalid "md5" parameter, BitBake returns an md5 mis-match error and displays the correct "md5" parameter value during the build. The correct parameter is also captured in the build log.

Tip

If the whole file contains only license text, you do not need to use the "beginline" and "endline" parameters.

3.4.2. Enabling Commercially Licensed Recipes

By default, the Yocto Project build system disables components that have commercial or other special licensing requirements. Such requirements are defined on a recipe-by-recipe basis through the `LICENSE_FLAGS` variable definition in the affected recipe. For instance, the `$HOME/poky/meta/recipes-multimedia/gstreamer/gst-plugins-ugly` recipe contains the following statement:

```
LICENSE_FLAGS = "commercial"
```

Here is a slightly more complicated example that contains both an explicit package name and version (after variable expansion):

```
LICENSE_FLAGS = "license_${PN}_${PV}"
```

In order for a component restricted by a `LICENSE_FLAGS` definition to be enabled and included in an image, it needs to have a matching entry in the global `LICENSE_FLAGS_WHITELIST` variable, which is a variable typically defined in your `local.conf` file. For example, to enable the `$HOME/poky/meta/recipes-multimedia/gstreamer/gst-plugins-ugly` package, you could add either the string `"commercial_gst-plugins-ugly"` or the more general string `"commercial"` to `LICENSE_FLAGS_WHITELIST`. See the "License Flag Matching" section for a full explanation of how `LICENSE_FLAGS` matching works. Here is the example:

```
LICENSE_FLAGS_WHITELIST = "commercial_gst-plugins-ugly"
```

Likewise, to additionally enable the package containing `LICENSE_FLAGS = "license_${PN}_${PV}"`, and assuming that the actual recipe name was `emgd_1.10.bb`, the following string would enable that package as well as the original `gst-plugins-ugly` package:

```
LICENSE_FLAGS_WHITELIST = "commercial_gst-plugins-ugly license_emgd_1.10"
```

As a convenience, you do not need to specify the complete license string in the whitelist for every package. you can use an abbreviated form, which consists of just the first portion or portions of the license string before the initial underscore character or characters. A partial string will match any license that contains the given string as the first portion of its license. For example, the following whitelist string will also match both of the packages previously mentioned as well as any other packages that have licenses starting with `"commercial"` or `"license"`.

```
LICENSE_FLAGS_WHITELIST = "commercial license"
```

3.4.2.1. License Flag Matching

The definition of 'matching' in reference to a recipe's `LICENSE_FLAGS` setting is simple. However, some things exist that you should know about in order to correctly and effectively use it.

Before a flag defined by a particular recipe is tested against the contents of the `LICENSE_FLAGS_WHITELIST` variable, the string `_${PN}` (with `PN` expanded of course) is appended to the flag, thus automatically making each `LICENSE_FLAGS` value recipe-specific. That string is then matched against the whitelist. So if you specify `LICENSE_FLAGS = "commercial"` in recipe `"foo"` for example, the string `"commercial_foo"` would normally be what is specified in the whitelist in order for it to match.

You can broaden the match by putting any `"_"`-separated beginning subset of a `LICENSE_FLAGS` flag in the whitelist, which will also match. For example, simply specifying `"commercial"` in the whitelist would match any expanded `LICENSE_FLAGS` definition starting with `"commercial"` such as `"commercial_foo"` and `"commercial_bar"`, which are the strings that would be automatically generated for hypothetical `"foo"` and `"bar"` recipes assuming those recipes had simply specified the following:

```
LICENSE_FLAGS = "commercial"
```

Broadening the match allows for a range of specificity for the items in the whitelist, from more general to perfectly specific. So you have the choice of exhaustively enumerating each license flag in the whitelist to allow only those specific recipes into the image, or of using a more general string to pick up anything matching just the first component or components of the specified string.

This scheme works even if the flag already has `_${PN}` appended - the extra `_${PN}` is redundant, but does not affect the outcome. For example, a license flag of `"commercial_1.2_foo"` would turn into `"commercial_1.2_foo_foo"` and would match both the general `"commercial"` and the specific `"commercial_1.2_foo"`, as expected. The flag would also match `"commercial_1.2_foo_foo"` and `"commercial_1.2"`, which does not make much sense regarding use in the whitelist.

For a versioned string, you could instead specify `"commercial_foo_1.2"`, which would turn into `"commercial_foo_1.2_foo"`. And, as expected, this flag allows you to pick up this package along with anything else `"commercial"` when you specify `"commercial"` in the whitelist. Or, the flag allows you to pick up this package along with anything `"commercial_foo"` regardless of version when you use `"commercial_foo"` in the whitelist. Finally, you can be completely specific about the package and version and specify `"commercial_foo_1.2"` package and version.

3.4.2.2. Other Variables Related to Commercial Licenses

Other helpful variables related to commercial license handling exist and are defined in the `$HOME/poky/meta/conf/distro/include/default-distrovars.inc` file:

```
COMMERCIAL_AUDIO_PLUGINS ?= ""
COMMERCIAL_VIDEO_PLUGINS ?= ""
COMMERCIAL_QT = ""
```

If you want to enable these components, you can do so by making sure you have the following statements in your `local.conf` configuration file:

```
COMMERCIAL_AUDIO_PLUGINS = "gst-plugins-ugly-mad \
    gst-plugins-ugly-mpegaudioparse"
COMMERCIAL_VIDEO_PLUGINS = "gst-plugins-ugly-mpeg2dec \
    gst-plugins-ugly-mpegstream gst-plugins-bad-mpegvideoparse"
COMMERCIAL_QT ?= "qmmmp"
LICENSE_FLAGS_WHITELIST = "commercial_gst-plugins-ugly commercial_gst-plugins-bad commercial_qt"
```

Of course, you could also create a matching whitelist for those components using the more general `"commercial"` in the whitelist, but that would also enable all the other packages with `LICENSE_FLAGS` containing `"commercial"`, which you may or may not want:

```
LICENSE_FLAGS_WHITELIST = "commercial"
```

Specifying audio and video plug-ins as part of the `COMMERCIAL_AUDIO_PLUGINS` and `COMMERCIAL_VIDEO_PLUGINS` statements or commercial qt components as part of the `COMMERCIAL_QT` statement (along with the enabling `LICENSE_FLAGS_WHITELIST`) includes the plug-ins or components into built images, thus adding support for media formats or components.

Chapter 4. Board Support Packages (BSP) - Developer's Guide

A Board Support Package (BSP) is a collection of information that defines how to support a particular hardware device, set of devices, or hardware platform. The BSP includes information about the hardware features present on the device and kernel configuration information along with any additional hardware drivers required. The BSP also lists any additional software components required in addition to a generic Linux software stack for both essential and optional platform features.

This chapter (or document if you are reading the BSP Developer's Guide) talks about BSP Layers, defines a structure for components so that BSPs follow a commonly understood layout, discusses how to customize a recipe for a BSP, addresses BSP licensing, and provides information that shows you how to create and manage a BSP Layer using two Yocto Project BSP Tools.

4.1. BSP Layers

The BSP consists of a file structure inside a base directory. Collectively, you can think of the base directory and the file structure as a BSP Layer. BSP Layers use the following naming convention:

```
meta-<bsp_name>
```

"bsp_name" is a placeholder for the machine or platform name.

The layer's base directory (meta-<bsp_name>) is the root of the BSP Layer. This root is what you add to the BBLAYERS [<http://www.yoctoproject.org/docs/1.2.2/poky-ref-manual/poky-ref-manual.html#var-BBLAYERS>] variable in the conf/bblayers.conf file found in the Yocto Project Build Directory [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#yocto-project-build-directory>]. Adding the root allows the Yocto Project build system to recognize the BSP definition and from it build an image. Here is an example:

```
BBLAYERS = " \
    /usr/local/src/yocto/meta \
    /usr/local/src/yocto/meta-yocto \
    /usr/local/src/yocto/meta-<bsp_name> \
    "
```

Some BSPs require additional layers on top of the BSP's root layer in order to be functional. For these cases, you also need to add those layers to the BBLAYERS variable in order to build the BSP. You must also specify in the "Dependencies" section of the BSP's README file any requirements for additional layers and, preferably, any build instructions that might be contained elsewhere in the README file.

Some layers function as a layer to hold other BSP layers. An example of this type of layer is the meta-intel layer. The meta-intel layer contains over 10 individual BSP layers.

For more detailed information on layers, see the "Understanding and Creating Layers" [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#understanding-and-creating-layers>] section of the Yocto Project Development Manual. You can also see the detailed examples in the appendices of The Yocto Project Development Manual [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html>].

4.2. Example Filesystem Layout

Providing a common form allows end-users to understand and become familiar with the layout. A common format also encourages standardization of software support of hardware.

The proposed form does have elements that are specific to the Yocto Project and OpenEmbedded build systems. It is intended that this information can be used by other systems besides Yocto Project

and OpenEmbedded and that it will be simple to extract information and convert it to other formats if required. Yocto Project, through its standard layers mechanism, can directly accept the format described as a layer. The BSP captures all the hardware-specific details in one place in a standard format, which is useful for any person wishing to use the hardware platform regardless of the build system they are using.

The BSP specification does not include a build system or other tools - it is concerned with the hardware-specific components only. At the end-distribution point, you can ship the BSP combined with a build system and other tools. However, it is important to maintain the distinction that these are separate components that happen to be combined in certain end products.

Before looking at the common form for the file structure inside a BSP Layer, you should be aware that some requirements do exist in order for a BSP to be considered compliant with the Yocto Project. For that list of requirements, see the "Released BSP Requirements" section.

Below is the common form for the file structure inside a BSP Layer. While you can use this basic form for the standard, realize that the actual structures for specific BSPs could differ.

```
meta-<bsp_name>/
meta-<bsp_name>/<bsp_license_file>
meta-<bsp_name>/README
meta-<bsp_name>/README.sources
meta-<bsp_name>/binary/<bootable_images>
meta-<bsp_name>/conf/layer.conf
meta-<bsp_name>/conf/machine/*.conf
meta-<bsp_name>/recipes-bsp/*
meta-<bsp_name>/recipes-core/*
meta-<bsp_name>/recipes-graphics/*
meta-<bsp_name>/recipes-kernel/linux/linux-yocto-<kernel_rev>.bbappend
```

Below is an example of the Crown Bay BSP:

```
meta-crownbay/COPYING.MIT
meta-crownbay/README
meta-crownbay/README.sources
meta-crownbay/binary/
meta-crownbay/conf/
meta-crownbay/conf/layer.conf
meta-crownbay/conf/machine/
meta-crownbay/conf/machine/crownbay.conf
meta-crownbay/conf/machine/crownbay-noemgd.conf
meta-crownbay/recipes-bsp/
meta-crownbay/recipes-bsp/formfactor/
meta-crownbay/recipes-bsp/formfactor/formfactor_0.0.bbappend
meta-crownbay/recipes-bsp/formfactor/formfactor/
meta-crownbay/recipes-bsp/formfactor/formfactor/crownbay/
meta-crownbay/recipes-bsp/formfactor/formfactor/crownbay/machconfig
meta-crownbay/recipes-bsp/formfactor/formfactor/crownbay-noemgd/
meta-crownbay/recipes-bsp/formfactor/formfactor/crownbay-noemgd/machconfig
meta-crownbay/recipes-core/
meta-crownbay/recipes-core/tasks/
meta-crownbay/recipes-core/tasks/task-core-tools-profile.bbappend
meta-crownbay/recipes-graphics/
meta-crownbay/recipes-graphics/xorg-xserver/
meta-crownbay/recipes-graphics/xorg-xserver/xserver-xf86-config_0.1.bbappend
meta-crownbay/recipes-graphics/xorg-xserver/xserver-xf86-config/
meta-crownbay/recipes-graphics/xorg-xserver/xserver-xf86-config/crownbay/
meta-crownbay/recipes-graphics/xorg-xserver/xserver-xf86-config/crownbay-xorg.conf
meta-crownbay/recipes-graphics/xorg-xserver/xserver-xf86-config/crownbay-noemgd/
meta-crownbay/recipes-graphics/xorg-xserver/xserver-xf86-config/crownbay-noemgd/xorg.conf
meta-crownbay/recipes-kernel/
meta-crownbay/recipes-kernel/linux/
meta-crownbay/recipes-kernel/linux/linux-yocto-rt_3.0.bbappend
```

```
meta-crownbay/recipes-kernel/linux/linux-yocto_2.6.37.bbappend  
meta-crownbay/recipes-kernel/linux/linux-yocto_3.0.bbappend
```

The following sections describe each part of the proposed BSP format.

4.2.1. License Files

You can find these files in the BSP Layer at:

```
meta-<bsp_name>/<bsp_license_file>
```

These optional files satisfy licensing requirements for the BSP. The type or types of files here can vary depending on the licensing requirements. For example, in the Crown Bay BSP all licensing requirements are handled with the COPYING.MIT file.

Licensing files can be MIT, BSD, GPLv*, and so forth. These files are recommended for the BSP but are optional and totally up to the BSP developer.

4.2.2. README File

You can find this file in the BSP Layer at:

```
meta-<bsp_name>/README
```

This file provides information on how to boot the live images that are optionally included in the binary/ directory. The README file also provides special information needed for building the image.

At a minimum, the README file must contain a list of dependencies, such as the names of any other layers on which the BSP depends and the name of the BSP maintainer with his or her contact information.

4.2.3. README.sources File

You can find this file in the BSP Layer at:

```
meta-<bsp_name>/README.sources
```

This file provides information on where to locate the BSP source files. For example, information provides where to find the sources that comprise the images shipped with the BSP. Information is also included to help you find the metadata used to generate the images that ship with the BSP.

4.2.4. Pre-built User Binaries

You can find these files in the BSP Layer at:

```
meta-<bsp_name>/binary/<bootable_images>
```

This optional area contains useful pre-built kernels and user-space filesystem images appropriate to the target system. This directory typically contains graphical (e.g. sato) and minimal live images when the BSP tarball has been created and made available in the Yocto Project [<http://www.yoctoproject.org>] website. You can use these kernels and images to get a system running and quickly get started on development tasks.

The exact types of binaries present are highly hardware-dependent. However, a README file should be present in the BSP Layer that explains how to use the kernels and images with the target hardware. If pre-built binaries are present, source code to meet licensing requirements must also exist in some form.

4.2.5. Layer Configuration File

You can find this file in the BSP Layer at:

```
meta-<bsp_name>/conf/layer.conf
```

The `conf/layer.conf` file identifies the file structure as a Yocto Project layer, identifies the contents of the layer, and contains information about how Yocto Project should use it. Generally, a standard boilerplate file such as the following works. In the following example, you would replace "bsp" and "_bsp" with the actual name of the BSP (i.e. <bsp_name> from the example template).

```
# We have a conf and classes directory, add to BBPATH
BBPATH := "${BBPATH}:${LAYERDIR}"

# We have a recipes directory, add to BBFILES
BBFILES := "${BBFILES} ${LAYERDIR}/recipes/*/*.bb \
           ${LAYERDIR}/recipes/*/*.bbappend"

BBFILE_COLLECTIONS += "bsp"
BBFILE_PATTERN_bsp := "^${LAYERDIR}/"
BBFILE_PRIORITY_bsp = "6"
```

To illustrate the string substitutions, here are the last three statements from the Crown Bay `conf/layer.conf` file:

```
BBFILE_COLLECTIONS += "crownbay"
BBFILE_PATTERN_crownbay := "^${LAYERDIR}/"
BBFILE_PRIORITY_crownbay = "6"
```

This file simply makes BitBake aware of the recipes and configuration directories. The file must exist so that the Yocto Project build system can recognize the BSP.

4.2.6. Hardware Configuration Options

You can find these files in the BSP Layer at:

```
meta-<bsp_name>/conf/machine/*.conf
```

The machine files bind together all the information contained elsewhere in the BSP into a format that the Yocto Project build system can understand. If the BSP supports multiple machines, multiple machine configuration files can be present. These filenames correspond to the values to which users have set the MACHINE [<http://www.yoctoproject.org/docs/1.2.2/poky-ref-manual/poky-ref-manual.html#var-MACHINE>] variable.

These files define things such as the kernel package to use (PREFERRED_PROVIDER [http://www.yoctoproject.org/docs/1.2.2/poky-ref-manual/poky-ref-manual.html#var-PREFERRED_PROVIDER] of virtual/kernel), the hardware drivers to include in different types of images, any special software components that are needed, any bootloader information, and also any special image format requirements.

Each BSP Layer requires at least one machine file. However, you can supply more than one file. For example, in the Crown Bay BSP shown earlier in this section, the `conf/machine` directory contains two configuration files: `crownbay.conf` and `crownbay-noemgd.conf`. The `crownbay.conf` file is used for the Crown Bay BSP that supports the Intel® Embedded Media and Graphics Driver (Intel® EMGD), while the `crownbay-noemgd.conf` file is used for the Crown Bay BSP that does not support the Intel® EMGD.

This `crownbay.conf` file could also include a hardware "tuning" file that is commonly used to define the package architecture and specify optimization flags, which are carefully chosen to give best performance on a given processor.

Tuning files are found in the `meta/conf/machine/include` directory of the Yocto Project Files [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#yocto-project-files>]. Tuning files can also reside in the BSP Layer itself. For example, the `ia32-base.inc` file resides in the `meta-intel` BSP Layer in `conf/machine/include`.

To use an include file, you simply include them in the machine configuration file. For example, the Crown Bay BSP `crownbay.conf` has the following statements:

```
include conf/machine/include/tune-atom.inc
include conf/machine/include/ia32-base.inc
```

4.2.7. Miscellaneous Recipe Files

You can find these files in the BSP Layer at:

```
meta-<bsp_name>/recipes-bsp/*
```

This optional directory contains miscellaneous recipe files for the BSP. Most notably would be the formfactor files. For example, in the Crown Bay BSP there is the `formfactor_0.0.bbappend` file, which is an append file used to augment the recipe that starts the build. Furthermore, there are machine-specific settings used during the build that are defined by the `machconfig` files. In the Crown Bay example, two `machconfig` files exist: one that supports the Intel® Embedded Media and Graphics Driver (Intel® EMGD) and one that does not:

```
meta-crownbay/recipes-bsp/formfactor/formfactor/crownbay/machconfig
meta-crownbay/recipes-bsp/formfactor/formfactor/crownbay-noemgd/machconfig
meta-crownbay/recipes-bsp/formfactor/formfactor_0.0.bbappend
```

Note

If a BSP does not have a formfactor entry, defaults are established according to the formfactor configuration file that is installed by the main formfactor recipe `meta/recipes-bsp/formfactor/formfactor_0.0.bb`, which is found in the Yocto Project Files [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#yocto-project-files>].

4.2.8. Core Recipe Files

You can find these files in the BSP Layer at:

```
meta-<bsp_name>/recipes-core/*
```

This directory contains recipe files that are almost always necessary to build a useful, working Linux image. Thus, the term "core" is used to group these recipes. For example, in the Crown Bay BSP there is the `task-core-tools-profile.bbappend` file, which is an append file used to recommend that the SystemTap [<http://sourceware.org/systemtap/wiki>] package be included as a package when the image is built.

4.2.9. Display Support Files

You can find these files in the BSP Layer at:

```
meta-<bsp_name>/recipes-graphics/*
```

This optional directory contains recipes for the BSP if it has special requirements for graphics support. All files that are needed for the BSP to support a display are kept here. For example, the Crown Bay BSP contains two versions of the `xorg.conf` file. The version in `crownbay` builds a BSP that supports the Intel® Embedded Media Graphics Driver (EMGD), while the version in `crownbay-noemgd` builds a BSP that supports Video Electronics Standards Association (VESA) graphics only:

```
meta-crownbay/recipes-graphics/xorg-xserver/xserver-xf86-config_0.1.bbappend
meta-crownbay/recipes-graphics/xorg-xserver/xserver-xf86-config/crownbay/xorg.conf
meta-crownbay/recipes-graphics/xorg-xserver/xserver-xf86-config/crownbay-noemgd/xorg.conf
```

4.2.10. Linux Kernel Configuration

You can find these files in the BSP Layer at:

```
meta-<bsp_name>/recipes-kernel/linux/linux-yocto_*.bbappend
```

These files append your specific changes to the kernel you are using.

For your BSP, you typically want to use an existing Yocto Project kernel found in the Yocto Project Files [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#yocto-project-files>] at `meta/recipes-kernel/linux`. You can append your specific changes to the kernel recipe by using a similarly named append file, which is located in the BSP Layer (e.g. the `meta-<bsp_name>/recipes-kernel/linux` directory).

Suppose the BSP uses the `linux-yocto_3.0.bb` kernel, which is the preferred kernel to use for developing a new BSP using the Yocto Project. In other words, you have selected the kernel in your `<bsp_name>.conf` file by adding the following statements:

```
PREFERRED_PROVIDER_virtual/kernel ?= "linux-yocto"
PREFERRED_VERSION_linux-yocto = "3.0%"
```

You would use the `linux-yocto_3.0.bbappend` file to append specific BSP settings to the kernel, thus configuring the kernel for your particular BSP.

As an example, look at the existing Crown Bay BSP. The append file used is:

```
meta-crownbay/recipes-kernel/linux/linux-yocto_3.0.bbappend
```

The following listing shows the file. Be aware that the actual commit ID strings in this example listing might be different than the actual strings in the file from the `meta-intel` Git source repository.

```
FILESEXTRAPATHS_prepend := "${THISDIR}/${PN}:"

COMPATIBLE_MACHINE_crownbay = "crownbay"
KMACHINE_crownbay = "yocto/standard/crownbay"
KERNEL_FEATURES_append_crownbay += " cfg/smp.scc"

COMPATIBLE_MACHINE_crownbay-noemgd = "crownbay-noemgd"
KMACHINE_crownbay-noemgd = "yocto/standard/crownbay"
KERNEL_FEATURES_append_crownbay-noemgd += " cfg/smp.scc"

SRCREV_machine_pn-linux-yocto_crownbay ?= "63c65842a3a74e4bd3128004ac29b5639f16433f"
SRCREV_meta_pn-linux-yocto_crownbay ?= "59314a3523e360796419d76d78c6f7d8c5ef2593"

SRCREV_machine_pn-linux-yocto_crownbay-noemgd ?= "63c65842a3a74e4bd3128004ac29b5639f16433f"
SRCREV_meta_pn-linux-yocto_crownbay-noemgd ?= "59314a3523e360796419d76d78c6f7d8c5ef2593"
```

This append file contains statements used to support the Crown Bay BSP for both Intel® EMGD and the VESA graphics. The build process, in this case, recognizes and uses only the statements that apply to the defined machine name - crownbay in this case. So, the applicable statements in the linux-yocto_3.0.bbappend file are follows:

```
FILESEXTRAPATHS_prepend := "${THISDIR}/${PN}:"

COMPATIBLE_MACHINE_crownbay = "crownbay"
KMACHINE_crownbay = "yocto/standard/crownbay"
KERNEL_FEATURES_append_crownbay += " cfg/smp.scc"

SRCREV_machine_pn-linux-yocto_crownbay ?= "63c65842a3a74e4bd3128004ac29b5639f16433f"
SRCREV_meta_pn-linux-yocto_crownbay ?= "59314a3523e360796419d76d78c6f7d8c5ef2593"
```

The append file defines crownbay as the compatible machine and defines the KMACHINE. The file also points to some configuration fragments to use by setting the KERNEL_FEATURES variable. The location for the configuration fragments is the kernel tree itself in the Yocto Project Build Directory [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#yocto-project-build-directory>] under linux/meta. Finally, the append file points to the specific commits in the Yocto Project Files [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#yocto-project-files>] Git repository and the meta Git repository branches to identify the exact kernel needed to build the Crown Bay BSP.

One thing missing in this particular BSP, which you will typically need when developing a BSP, is the kernel configuration file (.config) for your BSP. When developing a BSP, you probably have a kernel configuration file or a set of kernel configuration files that, when taken together, define the kernel configuration for your BSP. You can accomplish this definition by putting the configurations in a file or a set of files inside a directory located at the same level as your append file and having the same name as the kernel. With all these conditions met simply reference those files in a SRC_URI statement in the append file.

For example, suppose you had a set of configuration options in a file called myconfig. If you put that file inside a directory named /linux-yocto and then added a SRC_URI statement such as the following to the append file, those configuration options will be picked up and applied when the kernel is built.

```
SRC_URI += "file://myconfig"
```

As mentioned earlier, you can group related configurations into multiple files and name them all in the SRC_URI statement as well. For example, you could group separate configurations specifically for Ethernet and graphics into their own files and add those by using a SRC_URI statement like the following in your append file:

```
SRC_URI += "file://myconfig \
            file://eth.cfg \
            file://gfx.cfg"
```

The FILESEXTRAPATHS variable is in boilerplate form in the previous example in order to make it easy to do that. This variable must be in your layer or BitBake will not find the patches or configurations even if you have them in your SRC_URI. The FILESEXTRAPATHS variable enables the build process to find those configuration files.

Note

Other methods exist to accomplish grouping and defining configuration options. For example, if you are working with a local clone of the kernel repository, you could checkout the kernel's meta branch, make your changes, and then push the changes to the local bare clone of the kernel. The result is that you directly add configuration options to the Yocto kernel meta

branch for your BSP. The configuration options will likely end up in that location anyway if the BSP gets added to the Yocto Project. For an example showing how to change the BSP configuration, see the "Changing the BSP Configuration [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#changing-the-bsp-configuration>]" section in the Yocto Project Development Manual. For a better understanding of working with a local clone of the kernel repository and a local bare clone of the kernel, see the "Modifying the Kernel Source Code [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#modifying-the-kernel-source-code>]" section also in the Yocto Project Development Manual.

In general, however, the Yocto Project maintainers take care of moving the SRC_URI-specified configuration options to the kernel's meta branch. Not only is it easier for BSP developers to not have to worry about putting those configurations in the branch, but having the maintainers do it allows them to apply 'global' knowledge about the kinds of common configuration options multiple BSPs in the tree are typically using. This allows for promotion of common configurations into common features.

4.3. Requirements and Recommendations for Released BSPs

Certain requirements exist for a released BSP to be considered compliant with the Yocto Project. Additionally, a single recommendation also exists. This section describes the requirements and recommendation for released BSPs.

4.3.1. Released BSP Requirements

Before looking at BSP requirements, you should consider the following:

- The requirements here assume the BSP layer is a well-formed, "legal" layer that can be added to the Yocto Project. For guidelines on creating a Yocto Project layer that meets these base requirements, see the "BSP Layers" and the "Understanding and Creating Layers" [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#understanding-and-creating-layers>] in the Yocto Project Development Manual.
- The requirements in this section apply regardless of how you ultimately package a BSP. You should consult the packaging and distribution guidelines for your specific release process. For an example of packaging and distribution requirements, see the Third Party BSP Release Process [https://wiki.yoctoproject.org/wiki/Third_Party_BSP_Release_Process] wiki page.
- The requirements for the BSP as it is made available to a developer are completely independent of the released form of the BSP. For example, the BSP metadata can be contained within a Git repository and could have a directory structure completely different from what appears in the officially released BSP layer.
- It is not required that specific packages or package modifications exist in the BSP layer, beyond the requirements for general compliance with the Yocto Project. For example, no requirement exists dictating that a specific kernel or kernel version be used in a given BSP.

Following are the requirements for a released BSP that conforms to the Yocto Project:

- Layer Name: The BSP must have a layer name that follows the Yocto Project standards. For information on BSP layer names, see the "BSP Layers" section.
- File System Layout: When possible, use the same directory names in your BSP layer as listed in the `recipes.txt` file. In particular, you should place recipes (`.bb` files) and recipe modifications (`.bbappend` files) into `recipes-*` subdirectories by functional area as outlined in `recipes.txt`. If you cannot find a category in `recipes.txt` to fit a particular recipe, you can make up your own `recipe-*` subdirectory. You can find `recipes.txt` in the meta directory of the Yocto Project Files [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#yocto-project-files>], or in the OpenEmbedded Core Layer (`openembedded-core`) found at <http://git.openembedded.org/openembedded-core/tree/meta>.

Within any particular `recipes-*` category, the layout should match what is found in the OpenEmbedded Core Git repository (`openembedded-core`) or the Yocto Project Files (`poky`). In other words, make sure you place related files in appropriately related `recipes-*` subdirectories

specific to the recipe's function, or within a subdirectory containing a set of closely-related recipes. The recipes themselves should follow the general guidelines for recipes used in the Yocto Project found in the Yocto Recipe and Patch Style Guide [https://wiki.yoctoproject.org/wiki/Recipe_%26_Patch_Style_Guide].

- **License File:** You must include a license file in the meta-`<bsp_name>` directory. This license covers the BSP metadata as a whole. You must specify which license to use since there is no default license if one is not specified. See the `COPYING.MIT` [<http://git.yoctoproject.org/cgit.cgi/meta-intel/tree/meta-fishriver/COPYING.MIT>] file for the Fish River BSP in the meta-fishriver BSP layer as an example.
- **README File:** You must include a README file in the meta-`<bsp_name>` directory. See the README [<http://git.yoctoproject.org/cgit.cgi/meta-intel/tree/meta-fishriver/README>] file for the Fish River BSP in the meta-fishriver BSP layer as an example.

At a minimum, the README file should contain the following:

- A brief description about the hardware the BSP targets.
- A list of all the dependencies a on which a BSP layer depends. These dependencies are typically a list of required layers needed to build the BSP. However, the dependencies should also contain information regarding any other dependencies the BSP might have.
- Any required special licensing information. For example, this information includes information on special variables needed to satisfy a EULA, or instructions on information needed to build or distribute binaries built from the BSP metadata.
- The name and contact information for the BSP layer maintainer. This is the person to whom patches and questions should be sent.
- Instructions on how to build the BSP using the BSP layer.
- Instructions on how to boot the BSP build from the BSP layer.
- Instructions on how to boot the binary images contained in the `/binary` directory, if present.
- Information on any known bugs or issues that users should know about when either building or booting the BSP binaries.
- **README.sources File:** You must include a `README.sources` in the meta-`<bsp_name>` directory. This file specifies exactly where you can find the sources used to generate the binary images contained in the `/binary` directory, if present. See the `README.sources` [<http://git.yoctoproject.org/cgit.cgi/meta-intel/tree/meta-fishriver/README.sources>] file for the Fish River BSP in the meta-fishriver BSP layer as an example.
- **Layer Configuration File:** You must include a `conf/layer.conf` in the meta-`<bsp_name>` directory. This file identifies the meta-`<bsp_name>` BSP layer as a layer to the build system.
- **Machine Configuration File:** You must include a `conf/machine/<bsp_name>.conf` in the meta-`<bsp_name>` directory. This configuration file defines a machine target that can be built using the BSP layer. Multiple machine configuration files define variations of machine configurations that are supported by the BSP. If a BSP supports more multiple machine variations, you need to adequately describe each variation in the BSP README file. Do not use multiple machine configuration files to describe disparate hardware. Multiple machine configuration files should describe very similar targets. If you do have very different targets, you should create a separate BSP.

Note

It is completely possible for a developer to structure the working repository as a conglomeration of unrelated BSP files, and to possibly generate specifically targeted 'release' BSPs from that directory using scripts or some other mechanism. Such considerations are outside the scope of this document.

4.3.2. Released BSP Recommendations

Following are recommendations for a released BSP that conforms to the Yocto Project:

- **Bootable Images:** BSP releases can contain one or more bootable images. Including bootable images allows users to easily try out the BSP on their own hardware.

In some cases, it might not be convenient to include a bootable image. In this case, you might want to make two versions of the BSP available: one that contains binary images, and one that does not. The version that does not contain bootable images avoids unnecessary download times for users not interested in the images.

If you need to distribute a BSP and include bootable images or build kernel and filesystems meant to allow users to boot the BSP for evaluation purposes, you should put the images and artifacts within a binary/ subdirectory located in the meta-<bsp_name> directory.

Note

If you do include a bootable image as part of the BSP and the image was built by software covered by the GPL or other open source licenses, it is your responsibility to understand and meet all licensing requirements, which could include distribution of source files.

- Use a Yocto Linux Kernel: Kernel recipes in the BSP should be based on a Yocto Linux kernel. Basing your recipes on these kernels reduces the costs for maintaining the BSP and increases its scalability. See the Yocto Linux Kernel category in the Yocto Source Repositories [<http://git.yoctoproject.org/cgit.cgi>] for these kernels.

4.4. Customizing a Recipe for a BSP

If you plan on customizing a recipe for a particular BSP, you need to do the following:

- Include within the BSP layer a .bbappend file for the modified recipe.
- Place the BSP-specific file in the BSP's recipe .bbappend file path under a directory named after the machine.

To better understand this, consider an example that customizes a recipe by adding a BSP-specific configuration file named `interfaces` to the `netbase_4.47.bb` recipe for machine "xyz". Do the following:

1. Edit the `netbase_4.47.bbappend` file so that it contains the following:

```
FILESEXTRAPATHS_prepend := "${THISDIR}/files:"
PRINC := "${@int(PRINC) + 2}"
```

2. Create and place the new `interfaces` configuration file in the BSP's layer here:

```
meta-xyz/recipes-core/netbase/files/xyz/interfaces
```

4.5. BSP Licensing Considerations

In some cases, a BSP contains separately licensed Intellectual Property (IP) for a component or components. For these cases, you are required to accept the terms of a commercial or other type of license that requires some kind of explicit End User License Agreement (EULA). Once the license is accepted, the Yocto Project build system can then build and include the corresponding component in the final BSP image. If the BSP is available as a pre-built image, you can download the image after agreeing to the license or EULA.

You could find that some separately licensed components that are essential for normal operation of the system might not have an unencumbered (or free) substitute. Without these essential components, the system would be non-functional. Then again, you might find that other licensed components that are simply 'good-to-have' or purely elective do have an unencumbered, free replacement component that you can use rather than agreeing to the separately licensed component. Even for components essential to the system, you might find an unencumbered component that is not identical but will work as a less-capable version of the licensed version in the BSP recipe.

For cases where you can substitute a free component and still maintain the system's functionality, the Yocto Project website's BSP Download Page [<http://www.yoctoproject.org/download/>]

`all?keys=&download_type=1&download_version=]` makes available de-featured BSPs that are completely free of any IP encumbrances. For these cases, you can use the substitution directly and without any further licensing requirements. If present, these fully de-featured BSPs are named appropriately different as compared to the names of the respective encumbered BSPs. If available, these substitutions are your simplest and most preferred options. Use of these substitutions of course assumes the resulting functionality meets system requirements.

If however, a non-encumbered version is unavailable or it provides unsuitable functionality or quality, you can use an encumbered version.

A couple different methods exist within the Yocto Project build system to satisfy the licensing requirements for an encumbered BSP. The following list describes them in order of preference:

1. Use the `LICENSE_FLAGS` variable to define the Yocto Project recipes that have commercial or other types of specially-licensed packages: For each of those recipes, you can specify a matching license string in a `local.conf` variable named `LICENSE_FLAGS_WHITELIST`. Specifying the matching license string signifies that you agree to the license. Thus, the build system can build the corresponding recipe and include the component in the image. See the "Enabling Commercially Licensed Recipes [<http://www.yoctoproject.org/docs/1.2.2/poky-ref-manual/poky-ref-manual.html#enabling-commercially-licensed-recipes>]" section in the Yocto Project Reference Manual for details on how to use these variables.

If you build as you normally would, without specifying any recipes in the `LICENSE_FLAGS_WHITELIST`, the build stops and provides you with the list of recipes that you have tried to include in the image that need entries in the `LICENSE_FLAGS_WHITELIST`. Once you enter the appropriate license flags into the whitelist, restart the build to continue where it left off. During the build, the prompt will not appear again since you have satisfied the requirement.

Once the appropriate license flags are on the white list in the `LICENSE_FLAGS_WHITELIST` variable, you can build the encumbered image with no change at all to the normal build process.

2. Get a pre-built version of the BSP: You can get this type of BSP by visiting the Yocto Project website's Download [<http://www.yoctoproject.org/download>] page and clicking on "BSP Downloads". You can download BSP tarballs that contain proprietary components after agreeing to the licensing requirements of each of the individually encumbered packages as part of the download process. Obtaining the BSP this way allows you to access an encumbered image immediately after agreeing to the click-through license agreements presented by the website. Note that if you want to build the image yourself using the recipes contained within the BSP tarball, you will still need to create an appropriate `LICENSE_FLAGS_WHITELIST` to match the encumbered recipes in the BSP.

Note

Pre-compiled images are bundled with a time-limited kernel that runs for a predetermined amount of time (10 days) before it forces the system to reboot. This limitation is meant to discourage direct redistribution of the image. You must eventually rebuild the image if you want to remove this restriction.

4.6. Using the Yocto Project's BSP Tools

The Yocto Project includes a couple of tools that enable you to create a BSP layer from scratch and do basic configuration and maintenance of the kernel without ever looking at a Yocto Project metadata file. These tools are `yocto-bsp` and `yocto-kernel`, respectively.

The following sections describe the common location and help features as well as details for the `yocto-bsp` and `yocto-kernel` tools.

4.6.1. Common Features

Designed to have a command interface somewhat like Git [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#git>], each tool is structured as a set of sub-commands under a top-level command. The top-level command (`yocto-bsp` or `yocto-kernel`) itself does nothing but invoke or provide help on the sub-commands it supports.

Both tools reside in the `scripts/` subdirectory of the Yocto Project Files [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#yocto-project-files>]. Consequently, to use the scripts, you must source the environment just as you would when invoking a build:

```
$ source oe-init-build-env [build_dir]
```

The most immediately useful function is to get help on both tools. The built-in help system makes it easy to drill down at any time and view the syntax required for any specific command. Simply enter the name of the command, or the command along with help to display a list of the available sub-commands. Here is an example:

```
$ yocto-bsp
$ yocto-bsp help
```

Usage:

Create a customized Yocto BSP layer.

```
usage: yocto-bsp [--version] [--help] COMMAND [ARGS]
```

The most commonly used 'yocto-bsp' commands are:

```
create      Create a new Yocto BSP
list        List available values for options and BSP properties
```

See 'yocto-bsp help COMMAND' for more information on a specific command.

Options:

```
--version    show program's version number and exit
-h, --help   show this help message and exit
-D, --debug  output debug information
```

Similarly, entering just the name of a sub-command shows the detailed usage for that sub-command:

```
$ yocto-bsp create
```

Usage:

Create a new Yocto BSP

```
usage: yocto-bsp create <bsp-name> <karch> [-o <DIRNAME> | --outdir <DIRNAME>]
      [-i <JSON PROPERTY FILE> | --infile <JSON PROPERTY_FILE>]
```

This command creates a Yocto BSP based on the specified parameters. The new BSP will be a new Yocto BSP layer contained by default within the top-level directory specified as 'meta-bsp-name'. The -o option can be used to place the BSP layer in a directory with a different name and location.

...

For any sub-command, you can also use the word 'help' just before the sub-command to get more extensive documentation:

```
$ yocto-bsp help create
```

NAME

yocto-bsp create - Create a new Yocto BSP

SYNOPSIS

```
yocto-bsp create <bsp-name> <karch> [-o <DIRNAME> | --outdir <DIRNAME>]
      [-i <JSON PROPERTY FILE> | --infile <JSON PROPERTY_FILE>]
```

DESCRIPTION

This command creates a Yocto BSP based on the specified parameters. The new BSP will be a new Yocto BSP layer contained by default within the top-level directory specified as 'meta-bsp-name'. The -o option can be used to place the BSP layer in a directory with a different name and location.

The value of the 'karch' parameter determines the set of files that will be generated for the BSP, along with the specific set of 'properties' that will be used to fill out the BSP-specific portions of the BSP.

...

NOTE: Once created, you should add your new layer to your bblayers.conf file in order for it to be subsequently seen and modified by the yocto-kernel tool.

NOTE for x86- and x86_64-based BSPs: The generated BSP assumes the presence of the meta-intel layer, so you should also have a meta-intel layer present and added to your bblayers.conf as well.

Now that you know where these two commands reside and how to access information on them, you should find it relatively straightforward to discover the commands necessary to create a BSP and perform basic kernel maintenance on that BSP using the tools. The next sections provide a concrete starting point to expand on a few points that might not be immediately obvious or that could use further explanation.

4.6.2. Creating a new BSP Layer Using the yocto-bsp Script

The yocto-bsp script creates a new BSP layer for any architecture supported by the Yocto Project, as well as QEMU versions of the same. The default mode of the script's operation is to prompt you for information needed to generate the BSP layer. For the current set of BSPs, the script prompts you for various important parameters such as:

- which kernel to use
- which branch of that kernel to use (or re-use)
- whether or not to use X, and if so, which drivers to use
- whether to turn on SMP
- whether the BSP has a keyboard
- whether the BSP has a touchscreen
- any remaining configurable items associated with the BSP

You use the yocto-bsp create sub-command to create a new BSP layer. This command requires you to specify a particular architecture on which to base the BSP. Assuming you have sourced the environment, you can use the yocto-bsp list karch sub-command to list the architectures available for BSP creation as follows:

```
$ yocto-bsp list karch
Architectures available:
  arm
  powerpc
  i386
  mips
  x86_64
  qemu
```

The remainder of this section presents an example that uses `myarm` as the machine name and `qemu` as the machine architecture. Of the available architectures, `qemu` is the only architecture that causes the script to prompt you further for an actual architecture. In every other way, this architecture is representative of how creating a BSP for a 'real' machine would work. The reason the example uses this architecture is because it is an emulated architecture and can easily be followed without requiring actual hardware.

As the `yocto-bsp create` command runs, default values for the prompts appear in brackets. Pressing enter without supplying anything on the command line or pressing enter and providing an invalid response causes the script to accept the default value.

Following is the complete example:

```
$ yocto-bsp create myarm qemu
Which qemu architecture would you like to use? [default: x86]
  1) common 32-bit x86
  2) common 64-bit x86
  3) common 32-bit ARM
  4) common 32-bit PowerPC
  5) common 32-bit MIPS
3
Would you like to use the default (3.2) kernel? (Y/n)
Do you need a new machine branch for this BSP (the alternative is to re-use an existing branch)? (Y/n)
Getting branches from remote repo git://git.yoctoproject.org/linux-yocto-3.2...
Please choose a machine branch to base this BSP on => [default: standard/default/common-pc]
  1) base
  2) standard/base
  3) standard/default/arm-versatile-926ejs
  4) standard/default/base
  5) standard/default/beagleboard
  6) standard/default/cedartrailbsp (copy).xml
  7) standard/default/common-pc-64/base
  8) standard/default/common-pc-64/jasperforest
  9) standard/default/common-pc-64/romley
 10) standard/default/common-pc-64/sugarbay
 11) standard/default/common-pc/atom-pc
 12) standard/default/common-pc/base
 13) standard/default/crownbay
 14) standard/default/emenlow
 15) standard/default/fishriver
 16) standard/default/fri2
 17) standard/default/fsl-mpc8315e-rdb
 18) standard/default/mti-malta32-be
 19) standard/default/mti-malta32-le
 20) standard/default/preempt-rt
 21) standard/default/qemu-ppc32
 22) standard/default/routerstationpro
 23) standard/preempt-rt/base
 24) standard/preempt-rt/qemu-ppc32
 25) standard/preempt-rt/routerstationpro
 26) standard/tiny
3
Do you need SMP support? (Y/n)
Does your BSP have a touchscreen? (y/N)
Does your BSP have a keyboard? (Y/n)
New qemu BSP created in meta-myarm
```

Let's take a closer look at the example now:

1. For the `qemu` architecture, the script first prompts you for which emulated architecture to use. In the example, we use the `arm` architecture.

2. The script then prompts you for the kernel. The default kernel is 3.2 and is acceptable. So, the example accepts the default. If you enter 'n', the script prompts you to further enter the kernel you do want to use (e.g. 3.0, 3.2_preempt-rt, etc.).
3. Next, the script asks whether you would like to have a new branch created especially for your BSP in the local Linux Yocto Kernel [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#local-kernel-files>] Git repository . If not, then the script re-uses an existing branch.

In this example, the default (or 'yes') is accepted. Thus, a new branch is created for the BSP rather than using a common, shared branch. The new branch is the branch committed to for any patches you might later add. The reason a new branch is the default is that typically new BSPs do require BSP-specific patches. The tool thus assumes that most of time a new branch is required.

Note

In the current implementation, creation or re-use of a branch does not actually matter. The reason is because the generated BSPs assume that patches and configurations live in recipe-space, which is something that can be done with or without a dedicated branch. Generated BSPs, however, are different. This difference becomes significant once the tool's 'publish' functionality is implemented.

4. Regardless of which choice is made in the previous step, you are now given the opportunity to select a particular machine branch on which to base your new BSP-specific machine branch on (or to re-use if you had elected to not create a new branch). Because this example is generating an arm BSP, the example uses #3 at the prompt, which selects the arm-versatile branch.
5. The remainder of the prompts are routine. Defaults are accepted for each.
6. By default, the script creates the new BSP Layer in the Yocto Project Build Directory [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#yocto-project-build-directory>].

Once the BSP Layer is created, you must add it to your `bblayers.conf` file. Here is an example:

```
BBLAYERS = " \
    /usr/local/src/yocto/meta \
    /usr/local/src/yocto/meta-yocto \
    /usr/local/src/yocto/meta-myarm \
"
```

Adding the layer to this file allows the build system to build the BSP and the `yocto-kernel` tool to be able to find the layer and other metadata it needs on which to operate.

4.6.3. Managing Kernel Patches and Config Items with `yocto-kernel`

Assuming you have created a Yocto Project BSP Layer using `yocto-bsp` and you added it to your `BBLAYERS` [<http://www.yoctoproject.org/docs/1.2.2/poky-ref-manual/poky-ref-manual.html#var-BBLAYERS>] variable in the `bblayers.conf` file, you can now use the `yocto-kernel` script to add patches and configuration items to the BSP's kernel.

The `yocto-kernel` script allows you to add, remove, and list patches and kernel config settings to a Yocto Project BSP's kernel `.bbappend` file. All you need to do is use the appropriate sub-command. Recall that the easiest way to see exactly what sub-commands are available is to use the `yocto-kernel` built-in help as follows:

```
$ yocto-kernel
Usage:

Modify and list Yocto BSP kernel config items and patches.

usage: yocto-kernel [--version] [--help] COMMAND [ARGS]

The most commonly used 'yocto-kernel' commands are:
```

<code>config list</code>	List the modifiable set of bare kernel config options for a BSP
<code>config add</code>	Add or modify bare kernel config options for a BSP
<code>config rm</code>	Remove bare kernel config options from a BSP
<code>patch list</code>	List the patches associated with a BSP
<code>patch add</code>	Patch the Yocto kernel for a BSP
<code>patch rm</code>	Remove patches from a BSP

See 'yocto-kernel help COMMAND' for more information on a specific command.

The `yocto-kernel patch add` sub-command allows you to add a patch to a BSP. The following example adds two patches to the myarm BSP:

```
$ yocto-kernel patch add myarm ~/test.patch
Added patches:
    test.patch

$ yocto-kernel patch add myarm ~/yocto-testmod.patch
Added patches:
    yocto-testmod.patch
```

Note

Although the previous example adds patches one at a time, it is possible to add multiple patches at the same time.

You can verify patches have been added by using the `yocto-kernel patch list` sub-command. Here is an example:

```
$ yocto-kernel patch list myarm
The current set of machine-specific patches for myarm is:
    1) test.patch
    2) yocto-testmod.patch
```

You can also use the `yocto-kernel` script to remove a patch using the `yocto-kernel patch rm` sub-command. Here is an example:

```
$ yocto-kernel patch rm myarm
Specify the patches to remove:
    1) test.patch
    2) yocto-testmod.patch
1
Removed patches:
    test.patch
```

Again, using the `yocto-kernel patch list` sub-command, you can verify that the patch was in fact removed:

```
$ yocto-kernel patch list myarm
The current set of machine-specific patches for myarm is:
    1) yocto-testmod.patch
```

In a completely similar way, you can use the `yocto-kernel config add` sub-command to add one or more kernel config item settings to a BSP. The following commands add a couple of config items to the myarm BSP:

```
$ yocto-kernel config add myarm CONFIG_MISC_DEVICES=y
Added items:
```

```
CONFIG_MISC_DEVICES=y

$ yocto-kernel config add myarm KCONFIG_YOCTO_TESTMOD=y
Added items:
    CONFIG_YOCTO_TESTMOD=y
```

Note

Although the previous example adds config items one at a time, it is possible to add multiple config items at the same time.

You can list the config items now associated with the BSP. Doing so shows you the config items you added as well as others associated with the BSP:

```
$ yocto-kernel config list myarm
The current set of machine-specific kernel config items for myarm is:
    1) CONFIG_MISC_DEVICES=y
    2) CONFIG_YOCTO_TESTMOD=y
```

Finally, you can remove one or more config items using the `yocto-kernel config rm` sub-command in a manner completely analogous to `yocto-kernel patch rm`.

Chapter 5. Platform Development with the Yocto Project

5.1. Application Development Using the Yocto Project

The Yocto Project supports several methods of application development through which you can create user-space software designed to run on an embedded device that uses a Linux Yocto image developed with the Yocto Project. This flexibility allows you to choose the method that works best for you. This chapter describes each development method.

5.1.1. External Development Using the Meta-Toolchain

The Yocto Project provides toolchains that allow you to develop your application outside of the Yocto Project build system for specific hardware. These toolchains (called meta-toolchains) contain cross-development tools such as compilers, linkers, and debuggers that build your application for your target device. The Yocto Project also provides images that have toolchains for supported architectures included within the image. This allows you to compile, debug, or profile applications directly on the target device. See the "Reference: Images" appendix for a listing of the image types that Yocto Project supports.

Using the BitBake tool you can build a meta-toolchain or meta-toolchain-sdk target, which generates a tarball. Unpacking this tarball into the `/opt/poky` directory on your host produces a setup script (e.g. `/opt/poky/environment-setup-i586-poky-linux`) that you can source to initialize your build environment. Sourcing this script adds the compiler, QEMU scripts, QEMU binary, a special version of `pkgconfig` and other useful utilities to the `PATH` variable used by the Yocto Project build environment. Variables to assist `pkgconfig` and Autotools are also defined so that, for example, `configure` can find pre-generated test results for tests that need target hardware on which to run.

Using the toolchain with Autotool-enabled packages is straightforward - just pass the appropriate host option to `configure`. Following is an example:

```
$ ./configure --host=arm-poky-linux-gnueabi
```

For projects that are not Autotool-enabled, it is usually just a case of ensuring you point to and use the cross-toolchain. For example, the following two lines of code in a Makefile that builds your application specify to use the cross-compiler `arm-poky-linux-gnueabi-gcc` and linker `arm-poky-linux-gnueabi-ld`, which are part of the meta-toolchain you would have previously established:

```
CC=arm-poky-linux-gnueabi-gcc;  
LD=arm-poky-linux-gnueabi-ld;
```

5.1.2. External Development Using the Eclipse Plug-in

The current release of the Yocto Project supports the Eclipse IDE plug-in to make developing software easier for the application developer. The plug-in provides capability extensions to the graphical IDE to allow for cross compilation, deployment and execution of the application within a QEMU emulation session. Support of the Eclipse plug-in also allows for cross debugging and profiling. Additionally, the Eclipse plug-in provides a suite of tools that allows the developer to perform remote profiling, tracing, collection of power consumption data, collection of latency data and collection of performance data.

Note

The current release of the Yocto Project no longer supports the Anjuta plug-in. However, the Poky Anjuta Plug-in is available to download directly from the Poky Git repository located through the web interface at <http://git.yoctoproject.org> under IDE Plugins. The community is free to continue supporting it beyond the Yocto Project 0.9 Release.

To use the Eclipse plug-in you need the Eclipse Framework (Helios 3.6.1) along with other plug-ins installed into the Eclipse IDE. Once you have your environment setup you need to configure the Eclipse plug-in. For information on how to install and configure the Eclipse plug-in, see the "Working Within Eclipse [<http://www.yoctoproject.org/docs/1.2.2/adt-manual/adt-manual.html#adt-eclipse>]" chapter in the Yocto Project Application Development Toolkit (ADT) User's Guide.

5.1.3. External Development Using the QEMU Emulator

Running Poky QEMU images is covered in the "A Quick Test Run [<http://www.yoctoproject.org/docs/1.2.2/yocto-project-qs/yocto-project-qs.html#test-run>]" section of the Yocto Project Quick Start.

The QEMU images shipped with the Yocto Project contain complete toolchains native to their target architectures. This support allows you to develop applications within QEMU similar to the way you would using a normal host development system.

Speed can be an issue depending on the target and host architecture mix. For example, using the `qemux86` image in the emulator on an Intel-based 32-bit (x86) host machine is fast because the target and host architectures match. On the other hand, using the `qemuarm` image on the same Intel-based host can be slower. But, you still achieve faithful emulation of ARM-specific issues.

To speed things up, the QEMU images support using `distcc` to call a cross-compiler outside the emulated system. If you used `runqemu` to start QEMU, and `distccd` is present on the host system, any BitBake cross-compiling toolchain available from the build system is automatically used from within QEMU simply by calling `distcc`. You can accomplish this by defining the cross-compiler variable (e.g. `export CC="distcc"`). Alternatively, if a suitable SDK/toolchain is present in `/opt/poky` the toolchain is also automatically used.

Several mechanisms exist that let you connect to the system running on the QEMU emulator:

- QEMU provides a framebuffer interface that makes standard consoles available.
- Generally, headless embedded devices have a serial port. If so, you can configure the operating system of the running image to use that port to run a console. The connection uses standard IP networking.
- The QEMU images have a Dropbear secure shell (ssh) server that runs with the root password disabled. This allows you to use standard `ssh` and `scp` commands.
- The QEMU images also contain an embedded Network File System (NFS) server that exports the image's root filesystem. This allows you to make the filesystem available to the host.

5.1.4. Development Using Yocto Project Directly

Working directly with the Yocto Project is a fast and effective development technique. The idea is that you can directly edit files in a working directory (`WORKDIR`) or the source directory (`S`) and then force specific tasks to rerun in order to test the changes. An example session working on the `matchbox-desktop` package might look like this:

```
$ bitbake matchbox-desktop
$ sh
$ cd tmp/work/armv5te-poky-linux-gnueabi/matchbox-desktop-2.0+svnr1708-r0/
$ cd matchbox-desktop-2
$ vi src/main.c
.
.
[Make your changes]
.
.
$ exit
$ bitbake matchbox-desktop -c compile -f
$ bitbake matchbox-desktop
```

This example builds the `matchbox-desktop` package, creates a new terminal, changes into the work directory for the package, changes a file, exits out of the terminal, and then recompiles the package.

Instead of using `sh`, you can also use two different terminals. However, the risk of using two terminals is that a command like `unpack` could destroy your changes in the work directory. Consequently, you need to work carefully.

It is useful when making changes directly to the work directory files to do so using the Quilt tool as detailed in the "Using a Quilt Workflow [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#using-a-quilt-workflow>]" section in the Yocto Project Development Manual. Using Quilt, you can copy patches into the recipe directory and use the patches directly through use of the `SRC_URI` variable.

For a review of the skills used in this section, see the "BitBake" and "Running Specific Tasks" sections.

5.1.5. Development Within a Development Shell

When debugging certain commands or even when just editing packages, `devshell` can be a useful tool. Using `devshell` differs from the example shown in the previous section in that when you invoke `devshell` source files are extracted into your working directory and patches are applied. Then, a new terminal is opened and you are placed in the working directory. In the new terminal all the Yocto Project build-related environment variables are still defined so you can use commands such as `configure` and `make`. The commands execute just as if the Yocto Project build system were executing them. Consequently, working this way can be helpful when debugging a build or preparing software to be used with the Yocto Project build system.

Following is an example that uses `devshell` on a target named `matchbox-desktop`:

```
$ bitbake matchbox-desktop -c devshell
```

This command opens a terminal with a shell prompt within the Poky environment. The following occurs:

- The `PATH` variable includes the cross-toolchain.
- The `pkgconfig` variables find the correct `.pc` files.
- The `configure` command finds the Yocto Project site files as well as any other necessary files.

Within this environment, you can run `configure` or `compile` commands as if they were being run by the Yocto Project build system itself. As noted earlier, the working directory also automatically changes to the source directory (`S`).

When you are finished, you just exit the shell or close the terminal window.

The default shell used by `devshell` is `xterm`. For examples of available options, see the "UI/Interaction Configuration" section of the `meta/conf/bitbake.conf` configuration file in the Yocto Project files.

Because an external shell is launched rather than opening directly into the original terminal window, it allows easier interaction with BitBake's multiple threads as well as accomodates a future client/server split.

Note

It is worth remembering that when using `devshell` you need to use the full compiler name such as `arm-poky-linux-gnueabi-gcc` instead of just using `gcc`. The same applies to other applications such as `binutils`, `libtool` and so forth. BitBake sets up environment variables such as `CC` to assist applications, such as `make` to find the correct tools.

It is also worth noting that `devshell` still works over X11 forwarding and similar situations

5.1.6. Development Within Yocto Project for a Package that Uses an External SCM

If you're working on a recipe that pulls from an external Source Code Manager (SCM), it is possible to have the Yocto Project build system notice new changes added to the SCM and then build the package that depends on them using the latest version. This only works for SCMs from which it is possible

to get a sensible revision number for changes. Currently, you can do this with Apache Subversion (SVN), Git, and Bazaar (BZR) repositories.

To enable this behavior, simply add the following to the `local.conf` configuration file in the build directory of the Yocto Project files:

```
SRCREV_pn-<PN> = "${AUTOREV}"
```

where PN is the name of the package for which you want to enable automatic source revision updating.

5.2. Debugging With the GNU Project Debugger (GDB) Remotely

GDB allows you to examine running programs, which in turn help you to understand and fix problems. It also allows you to perform post-mortem style analysis of program crashes. GDB is available as a package within the Yocto Project and by default is installed in sdk images. See the "Reference: Images" appendix for a description of these images. You can find information on GDB at <http://sourceware.org/gdb/>.

Tip

For best results, install `-dbg` packages for the applications you are going to debug. Doing so makes available extra debug symbols that give you more meaningful output.

Sometimes, due to memory or disk space constraints, it is not possible to use GDB directly on the remote target to debug applications. These constraints arise because GDB needs to load the debugging information and the binaries of the process being debugged. Additionally, GDB needs to perform many computations to locate information such as function names, variable names and values, stack traces and so forth - even before starting the debugging process. These extra computations place more load on the target system and can alter the characteristics of the program being debugged.

To help get past the previously mentioned constraints, you can use Gdbserver. Gdbserver runs on the remote target and does not load any debugging information from the debugged process. Instead, a GDB instance processes the debugging information that is run on a remote computer - the host GDB. The host GDB then sends control commands to Gdbserver to make it stop or start the debugged program, as well as read or write memory regions of that debugged program. All the debugging information loaded and processed as well as all the heavy debugging is done by the host GDB. Offloading these processes gives the Gdbserver running on the target a chance to remain small and fast.

Because the host GDB is responsible for loading the debugging information and for doing the necessary processing to make actual debugging happen, the user has to make sure the host can access the unstripped binaries complete with their debugging information and also be sure the target is compiled with no optimizations. The host GDB must also have local access to all the libraries used by the debugged program. Because Gdbserver does not need any local debugging information, the binaries on the remote target can remain stripped. However, the binaries must also be compiled without optimization so they match the host's binaries.

To remain consistent with GDB documentation and terminology, the binary being debugged on the remote target machine is referred to as the "inferior" binary. For documentation on GDB see the GDB site [<http://sourceware.org/gdb/documentation/>].

5.2.1. Launching Gdbserver on the Target

First, make sure Gdbserver is installed on the target. If it is not, install the package `gdbserver`, which needs the `libthread-db1` package.

As an example, to launch Gdbserver on the target and make it ready to "debug" a program located at `/path/to/inferior`, connect to the target and launch:

```
$ gdbserver localhost:2345 /path/to/inferior
```

Gdbserver should now be listening on port 2345 for debugging commands coming from a remote GDB process that is running on the host computer. Communication between Gdbserver and the host GDB are done using TCP. To use other communication protocols, please refer to the Gdbserver documentation [<http://www.gnu.org/software/gdb/>].

5.2.2. Launching GDB on the Host Computer

Running GDB on the host computer takes a number of stages. This section describes those stages.

5.2.2.1. Building the Cross-GDB Package

A suitable GDB cross-binary is required that runs on your host computer but also knows about the ABI of the remote target. You can get this binary from the the Yocto Project meta-toolchain. Here is an example:

```
/usr/local/poky/eabi-glibc/arm/bin/arm-poky-linux-gnueabi-gdb
```

where `arm` is the target architecture and `linux-gnueabi` the target ABI.

Alternatively, the Yocto Project can build the `gdb-cross` binary. Here is an example:

```
$ bitbake gdb-cross
```

Once the binary is built, you can find it here:

```
tmp/sysroots/<host-arch>/usr/bin/<target-abi>-gdb
```

5.2.2.2. Making the Inferior Binaries Available

The inferior binary (complete with all debugging symbols) as well as any libraries (and their debugging symbols) on which the inferior binary depends need to be available. There are a number of ways you can make these available.

Perhaps the easiest way is to have an 'sdk' image that corresponds to the plain image installed on the device. In the case of `core-image-sato`, `core-image-sdk` would contain suitable symbols. Because the sdk images already have the debugging symbols installed, it is just a question of expanding the archive to some location and then informing GDB.

Alternatively, Yocto Project can build a custom directory of files for a specific debugging purpose by reusing its `tmp/rootfs` directory. This directory contains the contents of the last built image. This process assumes two things:

- The image running on the target was the last image to be built by the Yocto Project.
- The package (foo in the following example) that contains the inferior binary to be debugged has been built without optimization and has debugging information available.

The following steps show how to build the custom directory of files:

1. Install the package (foo in this case) to `tmp/rootfs`:

```
$ tmp/sysroots/i686-linux/usr/bin/opkg-cl -f \  
tmp/work/<target-abi>/core-image-sato-1.0-r0/temp/opkg.conf -o \  
tmp/rootfs/ update
```

2. Install the debugging information:

```
$ tmp/sysroots/i686-linux/usr/bin/opkg-cl -f \
tmp/work/<target-abi>/core-image-sato-1.0-r0/temp/opkg.conf \
-o tmp/rootfs install foo

$ tmp/sysroots/i686-linux/usr/bin/opkg-cl -f \
tmp/work/<target-abi>/core-image-sato-1.0-r0/temp/opkg.conf \
-o tmp/rootfs install foo-dbg
```

5.2.2.3. Launch the Host GDB

To launch the host GDB, you run the cross-gdb binary and provide the inferior binary as part of the command line. For example, the following command form continues with the example used in the previous section. This command form loads the foo binary as well as the debugging information:

```
$ <target-abi>-gdb rootfs/usr/bin/foo
```

Once the GDB prompt appears, you must instruct GDB to load all the libraries of the inferior binary from tmp/rootfs as follows:

```
$ set solib-absolute-prefix /path/to/tmp/rootfs
```

The pathname /path/to/tmp/rootfs must either be the absolute path to tmp/rootfs or the location at which binaries with debugging information reside.

At this point you can have GDB connect to the Gdbserver that is running on the remote target by using the following command form:

```
$ target remote remote-target-ip-address:2345
```

The remote-target-ip-address is the IP address of the remote target where the Gdbserver is running. Port 2345 is the port on which the GDBSERVER is running.

5.2.2.4. Using the Debugger

You can now proceed with debugging as normal - as if you were debugging on the local machine. For example, to instruct GDB to break in the "main" function and then continue with execution of the inferior binary use the following commands from within GDB:

```
(gdb) break main
(gdb) continue
```

For more information about using GDB, see the project's online documentation at <http://sourceware.org/gdb/download/onlinedocs/>.

5.3. Profiling with OProfile

OProfile [<http://oprofile.sourceforge.net/>] is a statistical profiler well suited for finding performance bottlenecks in both userspace software and in the kernel. This profiler provides answers to questions like "Which functions does my application spend the most time in when doing X?" Because the Yocto Project is well integrated with OProfile, it makes profiling applications on target hardware straightforward.

To use OProfile, you need an image that has OProfile installed. The easiest way to do this is with tools-profile in the IMAGE_FEATURES variable. You also need debugging symbols to be available on the system where the analysis takes place. You can gain access to the symbols by using dbg-pkgs in the IMAGE_FEATURES variable or by installing the appropriate -dbg packages.

For successful call graph analysis, the binaries must preserve the frame pointer register and should also be compiled with the `-fno-omit-framepointer` flag. In the Yocto Project you can achieve this by setting the `SELECTED_OPTIMIZATION` variable to `-fexpensive-optimizations -fno-omit-framepointer -frename-registers -O2`. You can also achieve it by setting the `DEBUG_BUILD` variable to "1" in the `local.conf` configuration file. If you use the `DEBUG_BUILD` variable you will also add extra debug information that can make the debug packages large.

5.3.1. Profiling on the Target

Using OProfile you can perform all the profiling work on the target device. A simple OProfile session might look like the following:

```
# opcontrol --reset
# opcontrol --start --separate=lib --no-vmlinux -c 5
.
.
[do whatever is being profiled]
.
.
# opcontrol --stop
$ opreport -cl
```

In this example, the `reset` command clears any previously profiled data. The next command starts OProfile. The options used when starting the profiler separate dynamic library data within applications, disable kernel profiling, and enable callgraphing up to five levels deep.

Note

To profile the kernel, you would specify the `--vmlinux=/path/to/vmlinux` option. The `vmlinux` file is usually in the Yocto Project file's `/boot/` directory and must match the running kernel.

After you perform your profiling tasks, the next command stops the profiler. After that, you can view results with the `opreport` command with options to see the separate library symbols and callgraph information.

Callgraphing logs information about time spent in functions and about a function's calling function (parent) and called functions (children). The higher the callgraphing depth, the more accurate the results. However, higher depths also increase the logging overhead. Consequently, you should take care when setting the callgraphing depth.

Note

On ARM, binaries need to have the frame pointer enabled for callgraphing to work. To accomplish this use the `-fno-omit-framepointer` option with `gcc`.

For more information on using OProfile, see the OProfile online documentation at <http://oprofile.sourceforge.net/docs/>.

5.3.2. Using OProfileUI

A graphical user interface for OProfile is also available. You can download and build this interface from the Yocto Project at <http://git.yoctoproject.org/cgit.cgi/oprofileui/>. If the "tools-profile" image feature is selected, all necessary binaries are installed onto the target device for OProfileUI interaction.

Even though the Yocto Project usually includes all needed patches on the target device, you might find you need other OProfile patches for recent OProfileUI features. If so, see the OProfileUI README [<http://git.yoctoproject.org/cgit.cgi/oprofileui/tree/README>] for the most recent information.

5.3.2.1. Online Mode

Using OProfile in online mode assumes a working network connection with the target hardware. With this connection, you just need to run "oprofile-server" on the device. By default, OProfile listens on port 4224.

Note

You can change the port using the `--port` command-line option.

The client program is called `oprofile-viewer` and its UI is relatively straightforward. You access key functionality through the buttons on the toolbar, which are duplicated in the menus. Here are the buttons:

- **Connect:** Connects to the remote host. You can also supply the IP address or hostname.
- **Disconnect:** Disconnects from the target.
- **Start:** Starts profiling on the device.
- **Stop:** Stops profiling on the device and downloads the data to the local host. Stopping the profiler generates the profile and displays it in the viewer.
- **Download:** Downloads the data from the target and generates the profile, which appears in the viewer.
- **Reset:** Resets the sample data on the device. Resetting the data removes sample information collected from previous sampling runs. Be sure you reset the data if you do not want to include old sample information.
- **Save:** Saves the data downloaded from the target to another directory for later examination.
- **Open:** Loads previously saved data.

The client downloads the complete 'profile archive' from the target to the host for processing. This archive is a directory that contains the sample data, the object files, and the debug information for the object files. The archive is then converted using the `oparchconv` script, which is included in this distribution. The script uses `opimport` to convert the archive from the target to something that can be processed on the host.

Downloaded archives reside in the Yocto Project's build directory in `/tmp` and are cleared up when they are no longer in use.

If you wish to perform kernel profiling, you need to be sure a `vmlinux` file that matches the running kernel is available. In the Yocto Project, that file is usually located in `/boot/vmlinux-KERNELVERSION`, where `KERNEL-version` is the version of the kernel. The Yocto Project generates separate `vmlinux` packages for each kernel it builds. Thus, it should just be a question of making sure a matching package is installed (e.g. `opkg install kernel-vmlinux`). The files are automatically installed into development and profiling images alongside `OProfile`. A configuration option exists within the `OProfileUI` settings page that you can use to enter the location of the `vmlinux` file.

Waiting for debug symbols to transfer from the device can be slow, and it is not always necessary to actually have them on the device for `OProfile` use. All that is needed is a copy of the filesystem with the debug symbols present on the viewer system. The "Launching GDB on the Host Computer" section covers how to create such a directory with the Yocto Project and how to use the `OProfileUI` Settings dialog to specify the location. If you specify the directory, it will be used when the file checksums match those on the system you are profiling.

5.3.2.2. Offline Mode

If network access to the target is unavailable, you can generate an archive for processing in `oprofile-viewer` as follows:

```
# opcontrol --reset
# opcontrol --start --separate=lib --no-vmlinux -c 5
.
.
[do whatever is being profiled]
.
.
# opcontrol --stop
# oparchive -o my_archive
```

In the above example, `my_archive` is the name of the archive directory where you would like the profile archive to be kept. After the directory is created, you can copy it to another host and load it using `oprofile-viewer` open functionality. If necessary, the archive is converted.

Appendix A. Reference: Directory Structure

The Yocto Project consists of several components. Understanding them and knowing where they are located is key to using the Yocto Project well. This appendix describes the Yocto Project file's directory structure and gives information about the various files and directories.

For information on how to establish the Yocto Project files on your local development system, see the "Getting Set Up [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#getting-setup>]" section in the Yocto Project Development Manual.

A.1. Top level core components

A.1.1. **bitbake/**

The Yocto Project includes a copy of BitBake for ease of use. The copy usually matches the current stable BitBake release from the BitBake project. BitBake, a metadata interpreter, reads the Yocto Project metadata and runs the tasks defined by that data. Failures are usually from the metadata and not from BitBake itself. Consequently, most users do not need to worry about BitBake. The `bitbake/` directory is placed into the shell's PATH environment variable by the `oe-init-build-env` script.

For more information on BitBake, see the BitBake on-line manual at <http://bitbake.berlios.de/manual/>.

A.1.2. **build/**

This directory contains user configuration files and the output generated by the Yocto Project in its standard configuration where the source tree is combined with the output. The build directory is created initially when you source the Yocto Project environment setup script `oe-init-build-env`.

It is also possible to place output and configuration files in a directory separate from the Yocto Project files by providing a directory name when you source the setup script. For information on separating output from the Yocto Project files, see `oe-init-build-env`.

A.1.3. **documentation**

This directory holds the source for the Yocto Project documentation as well as templates and tools that allow you to generate PDF and HTML versions of the manuals. Each manual is contained in a sub-folder. For example, the files for this manual reside in `poky-ref-manual`.

A.1.4. **meta/**

This directory contains the Yocto Project core metadata. The directory holds machine definitions, the Yocto Project distribution, and the packages that make up a given system.

A.1.5. **meta-demoapps/**

This directory contains recipes for applications and demos that are not part of the Yocto Project core.

A.1.6. **meta-rt/**

This directory contains recipes for real-time kernels.

A.1.7. **meta-skeleton/**

This directory contains template recipes for BSP and kernel development.

A.1.8. **scripts/**

This directory contains various integration scripts that implement extra functionality in the Yocto Project environment (e.g. QEMU scripts). The `oe-init-build-env` script appends this directory to the shell's PATH environment variable.

The `scripts` directory has useful scripts that assist contributing back to the Yocto Project, such as `create_pull_request` and `send_pull_request`.

A.1.9. **oe-init-build-env**

This script sets up the Yocto Project build environment. Running this script with the `source` command in a shell makes changes to `PATH` and sets other core BitBake variables based on the current working directory. You need to run this script before running BitBake commands. The script uses other scripts within the `scripts` directory to do the bulk of the work.

By default, running this script without a build directory argument creates the build directory. If you provide a build directory argument when you source the script, you direct the Yocto Project to create a build directory of your choice. For example, the following command creates a build directory named `mybuilds` that is outside of the Yocto Project files:

```
$ source oe-init-build-env ~/mybuilds
```

A.1.10. **LICENSE, README, and README.hardware**

These files are standard top-level files.

A.2. The Build Directory **-build/**

A.2.1. **build/pseudodone**

This tag file indicates that the initial pseudo binary was created. The file is built the first time BitBake is invoked.

A.2.2. **build/conf/local.conf**

This file contains all the local user configuration of the Yocto Project. If there is no `local.conf` present, it is created from `local.conf.sample`. The `local.conf` file contains documentation on the various configuration options. Any variable set here overrides any variable set elsewhere within the Yocto Project unless that variable is hard-coded within the Yocto Project (e.g. by using `'='` instead of `'?='`). Some variables are hard-coded for various reasons but these variables are relatively rare.

Edit this file to set the `MACHINE` for which you want to build, which package types you wish to use (`PACKAGE_CLASSES`), or where you want to download files (`DL_DIR`).

A.2.3. **build/conf/bblayers.conf**

This file defines layers, which is a directory tree, traversed (or walked) by BitBake. If `bblayers.conf` is not present, it is created from `bblayers.conf.sample` when you source the environment setup script.

A.2.4. **build/conf/sanity_info**

This file is created during the build to indicate the state of the sanity checks.

A.2.5. **build/downloads/**

This directory is used for the upstream source tarballs. The directory can be reused by multiple builds or moved to another location. You can control the location of this directory through the `DL_DIR` variable.

A.2.6. **build/sstate-cache/**

This directory is used for the shared state cache. The directory can be reused by multiple builds or moved to another location. You can control the location of this directory through the `SSTATE_DIR` variable.

A.2.7. **build/tmp/**

This directory receives all the Yocto Project output. BitBake creates this directory if it does not exist. As a last resort, to clean the Yocto Project and start a build from scratch (other than downloads), you can remove everything in this directory or get rid of the directory completely. If you do, you should also completely remove the `build/sstate-cache` directory as well.

A.2.8. **build/tmp/buildstats/**

This directory stores the build statistics.

A.2.9. **build/tmp/cache/**

When BitBake parses the metadata, it creates a cache file of the result that can be used when subsequently running commands. These results are stored here on a per-machine basis.

A.2.10. **build/tmp/deploy/**

This directory contains any 'end result' output from the Yocto Project build process.

A.2.11. **build/tmp/deploy/deb/**

This directory receives any `.deb` packages produced by the Yocto Project. The packages are sorted into feeds for different architecture types.

A.2.12. **build/tmp/deploy/rpm/**

This directory receives any `.rpm` packages produced by the Yocto Project. The packages are sorted into feeds for different architecture types.

A.2.13. **build/tmp/deploy/images/**

This directory receives complete filesystem images. If you want to flash the resulting image from a build onto a device, look here for the image.

Note, you should not remove any files from this directory by hand in an attempt to rebuild an image. If you want to clean out the cache, re-run the build using the following BitBake command:

```
$ bitbake -c cleanall <target>
```

A.2.14. **build/tmp/deploy/ipk/**

This directory receives `.ipk` packages produced by the Yocto Project.

A.2.15. **build/tmp/sysroots/**

This directory contains shared header files and libraries as well as other shared data. Packages that need to share output with other packages do so within this directory. The directory is subdivided by architecture so multiple builds can run within the one build directory.

A.2.16. **build/tmp/stamps/**

This directory holds information that that BitBake uses for accounting purposes to track what tasks have run and when they have run. The directory is sub-divided by architecture. The files in the directory are empty of data. However, BitBake uses the filenames and timestamps for tracking purposes.

A.2.17. **build/tmp/log/**

This directory contains general logs that are not otherwise placed using the package's `WORKDIR`. Examples of logs are the output from the `check_pkg` or `distro_check` tasks.

A.2.18. **build/tmp/pkgdata/**

This directory contains intermediate packaging data that is used later in the packaging process. For more information, see the "Packaging - package*.bbclass" section.

A.2.19. **build/tmp/work/**

This directory contains architecture-specific work sub-directories for packages built by BitBake. All tasks execute from a work directory. For example, the source for a particular package is unpacked, patched, configured and compiled all within its own work directory. Within the work directory, organization is based on the package group for which the source is being compiled.

It is worth considering the structure of a typical work directory. As an example, consider the linux-yocto kernel 3.0 on the machine qemu86 built within the Yocto Project. For this package, a work directory of tmp/work/qemu86-poky-linux/linux-yocto-3.0+git1+<...>, referred to as WORKDIR, is created. Within this directory, the source is unpacked to linux-qemu86-standard-build and then patched by Quilt (see the "Modifying Package Source Code with Quilt [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#using-a-quilt-workflow>]" section in the Yocto Project Development Manual. Within the linux-qemu86-standard-build directory, standard Quilt directories linux-3.0/patches and linux-3.0/.pc are created, and standard Quilt commands can be used.

There are other directories generated within WORKDIR. The most important directory is WORKDIR/temp/, which has log files for each task (log.do_*.pid) and contains the scripts BitBake runs for each task (run.do_*.pid). The WORKDIR/image/ directory is where "make install" places its output that is then split into sub-packages within WORKDIR/packages-split/.

A.3. The Metadata -meta/

As mentioned previously, metadata is the core of the Yocto Project. Metadata has several important subdivisions:

A.3.1. **meta/classes/**

This directory contains the *.bbclass files. Class files are used to abstract common code so it can be reused by multiple packages. Every package inherits the base.bbclass file. Examples of other important classes are autotools.bbclass, which in theory allows any Autotool-enabled package to work with the Yocto Project with minimal effort. Another example is kernel.bbclass that contains common code and functions for working with the Linux kernel. Functions like image generation or packaging also have their specific class files such as image.bbclass, rootfs_*.bbclass and package*.bbclass.

A.3.2. **meta/conf/**

This directory contains the core set of configuration files that start from bitbake.conf and from which all other configuration files are included. See the include statements at the end of the file and you will note that even local.conf is loaded from there. While bitbake.conf sets up the defaults, you can often override these by using the (local.conf) file, machine file or the distribution configuration file.

A.3.3. **meta/conf/machine/**

This directory contains all the machine configuration files. If you set MACHINE="qemu86", Yocto Project looks for a qemu86.conf file in this directory. The include directory contains various data common to multiple machines. If you want to add support for a new machine to the Yocto Project, look in this directory.

A.3.4. **meta/conf/distro/**

Any distribution-specific configuration is controlled from this directory. The Yocto Project only contains the Yocto Project distribution so defaultsetup.conf is the main file here. This directory includes the versions and the SRCDATE definitions for applications that are configured here. An example of an

alternative configuration is `poky-bleeding.conf` although this file mainly inherits its configuration from the Yocto Project itself.

A.3.5. meta/recipes-bsp/

This directory contains anything linking to specific hardware or hardware configuration information such as "u-boot" and "grub".

A.3.6. meta/recipes-connectivity/

This directory contains libraries and applications related to communication with other devices.

A.3.7. meta/recipes-core/

This directory contains what is needed to build a basic working Linux image including commonly used dependencies.

A.3.8. meta/recipes-devtools/

This directory contains tools that are primarily used by the build system. The tools, however, can also be used on targets.

A.3.9. meta/recipes-extended/

This directory contains non-essential applications that add features compared to the alternatives in core. You might need this directory for full tool functionality or for Linux Standard Base (LSB) compliance.

A.3.10. meta/recipes-gnome/

This directory contains all things related to the GTK+ application framework.

A.3.11. meta/recipes-graphics/

This directory contains X and other graphically related system libraries

A.3.12. meta/recipes-kernel/

This directory contains the kernel and generic applications and libraries that have strong kernel dependencies.

A.3.13. meta/recipes-multimedia/

This directory contains codecs and support utilities for audio, images and video.

A.3.14. meta/recipes-qt/

This directory contains all things related to the Qt application framework.

A.3.15. meta/recipes-sato/

This directory contains the Sato demo/reference UI/UX and its associated applications and configuration data.

A.3.16. meta/recipes-support/

This directory contains recipes that used by other recipes, but that are not directly included in images (i.e. dependencies of other recipes).

A.3.17. meta/site/

This directory contains a list of cached results for various architectures. Because certain "autoconf" test results cannot be determined when cross-compiling due to the tests not able to run on a live system, the information in this directory is passed to "autoconf" for the various architectures.

A.3.18. meta/recipes.txt/

This file is a description of the contents of recipes-.*.

Appendix B. Reference: BitBake

BitBake is a program written in Python that interprets the metadata that makes up the Yocto Project. At some point, developers wonder what actually happens when you enter:

```
$ bitbake core-image-sato
```

This appendix provides an overview of what happens behind the scenes from BitBake's perspective.

Note

BitBake strives to be a generic "task" executor that is capable of handling complex dependency relationships. As such, it has no real knowledge of what the tasks being executed actually do. BitBake just considers a list of tasks with dependencies and handles metadata that consists of variables in a certain format that get passed to the tasks.

B.1. Parsing

BitBake parses configuration files, classes, and .bb files.

The first thing BitBake does is look for the `bitbake.conf` file. The Yocto Project keeps this file in the Yocto Project file's `meta/conf/` directory. BitBake finds it by examining its `BBPATH` environment variable and looking for the `meta/conf/` directory.

In the Yocto Project, `bitbake.conf` lists other configuration files to include from a `conf/` directory below the directories listed in `BBPATH`. In general, the most important configuration file from a user's perspective is `local.conf`, which contains a user's customized settings for the Yocto Project build environment. Other notable configuration files are the distribution configuration file (set by the `DISTRO` variable) and the machine configuration file (set by the `MACHINE` variable). The `DISTRO` and `MACHINE` BitBake environment variables are both usually set in the `local.conf` file. Valid distribution configuration files are available in the `meta/conf/distro/` directory and valid machine configuration files in the `meta/conf/machine/` directory. Within the `meta/conf/machine/include/` directory are various `tune-*.inc` configuration files that provide common "tuning" settings specific to and shared between particular architectures and machines.

After the parsing of the configuration files, some standard classes are included. The `base.bbclass` file is always included. Other classes that are specified in the configuration using the `INHERIT` variable are also included. Class files are searched for in a `classes` subdirectory under the paths in `BBPATH` in the same way as configuration files.

After classes are included, the variable `BBFILES` is set, usually in `local.conf`, and defines the list of places to search for .bb files. By default, the `BBFILES` variable specifies the `meta/recipes-*/` directory within Poky. Adding extra content to `BBFILES` is best achieved through the use of BitBake layers as described in the "Understanding and Creating Layers [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#understanding-and-creating-layers>]" section of the Yocto Project Development Manual.

BitBake parses each .bb file in `BBFILES` and stores the values of various variables. In summary, for each .bb file the configuration plus the base class of variables are set, followed by the data in the .bb file itself, followed by any `inherit` commands that .bb file might contain.

Because parsing .bb files is a time consuming process, a cache is kept to speed up subsequent parsing. This cache is invalid if the timestamp of the .bb file itself changes, or if the timestamps of any of the include, configuration or class files the .bb file depends on changes.

B.2. Preferences and Providers

Once all the .bb files have been parsed, BitBake starts to build the target (`core-image-sato` in the previous section's example) and looks for providers of that target. Once a provider is selected, BitBake resolves all the dependencies for the target. In the case of `core-image-sato`, it would lead to `task-base.bb`, which in turn leads to packages like `Contacts`, `Dates` and `BusyBox`. These packages in turn depend on `eglibc` and the toolchain.

Sometimes a target might have multiple providers. A common example is "virtual/kernel", which is provided by each kernel package. Each machine often selects the best kernel provider by using a line similar to the following in the machine configuration file:

```
PREFERRED_PROVIDER_virtual/kernel = "linux-yocto"
```

The default `PREFERRED_PROVIDER` is the provider with the same name as the target.

Understanding how providers are chosen is made complicated by the fact that multiple versions might exist. BitBake defaults to the highest version of a provider. Version comparisons are made using the same method as Debian. You can use the `PREFERRED_VERSION` variable to specify a particular version (usually in the distro configuration). You can influence the order by using the `DEFAULT_PREFERENCE` variable. By default, files have a preference of "0". Setting the `DEFAULT_PREFERENCE` to "-1" makes the package unlikely to be used unless it is explicitly referenced. Setting the `DEFAULT_PREFERENCE` to "1" makes it likely the package is used. `PREFERRED_VERSION` overrides any `DEFAULT_PREFERENCE` setting. `DEFAULT_PREFERENCE` is often used to mark newer and more experimental package versions until they have undergone sufficient testing to be considered stable.

In summary, BitBake has created a list of providers, which is prioritized, for each target.

B.3. Dependencies

Each target BitBake builds consists of multiple tasks such as `fetch`, `unpack`, `patch`, `configure`, and `compile`. For best performance on multi-core systems, BitBake considers each task as an independent entity with its own set of dependencies.

Dependencies are defined through several variables. You can find information about variables BitBake uses in the BitBake manual [<http://bitbake.berlios.de/manual/>]. At a basic level, it is sufficient to know that BitBake uses the `DEPENDS` and `RDEPENDS` variables when calculating dependencies.

B.4. The Task List

Based on the generated list of providers and the dependency information, BitBake can now calculate exactly what tasks it needs to run and in what order it needs to run them. The build now starts with BitBake forking off threads up to the limit set in the `BB_NUMBER_THREADS` variable. BitBake continues to fork threads as long as there are tasks ready to run, those tasks have all their dependencies met, and the thread threshold has not been exceeded.

It is worth noting that you can greatly speed up the build time by properly setting the `BB_NUMBER_THREADS` variable. See the "Building an Image [<http://www.yoctoproject.org/docs/1.2.2/yocto-project-qs/yocto-project-qs.html#building-image>]" section in the Yocto Project Quick Start for more information.

As each task completes, a timestamp is written to the directory specified by the `STAMP` variable (usually `build/tmp/stamps/*/`). On subsequent runs, BitBake looks at the `/build/tmp/stamps` directory and does not rerun tasks that are already completed unless a timestamp is found to be invalid. Currently, invalid timestamps are only considered on a per `.bb` file basis. So, for example, if the `configure` stamp has a timestamp greater than the `compile` timestamp for a given target, then the `compile` task would rerun. Running the `compile` task again, however, has no effect on other providers that depend on that target. This behavior could change or become configurable in future versions of BitBake.

Note

Some tasks are marked as "nostamp" tasks. No timestamp file is created when these tasks are run. Consequently, "nostamp" tasks are always rerun.

B.5. Running a Task

Tasks can either be a shell task or a Python task. For shell tasks, BitBake writes a shell script to `${WORKDIR}/temp/run.do_taskname.pid` and then executes the script. The generated shell script contains all the exported variables, and the shell functions with all variables expanded. Output from

the shell script goes to the file `${WORKDIR}/temp/log.do_taskname.pid`. Looking at the expanded shell functions in the run file and the output in the log files is a useful debugging technique.

For Python tasks, BitBake executes the task internally and logs information to the controlling terminal. Future versions of BitBake will write the functions to files similar to the way shell tasks are handled. Logging will be handled in way similar to shell tasks as well.

Once all the tasks have been completed BitBake exits.

When running a task, BitBake tightly controls the execution environment of the build tasks to make sure unwanted contamination from the build machine cannot influence the build. Consequently, if you do want something to get passed into the build task's environment, you must take a few steps:

1. Tell BitBake to load what you want from the environment into the data store. You can do so through the `BB_ENV_WHITELIST` variable. For example, assume you want to prevent the build system from accessing your `$HOME/.ccache` directory. The following command tells BitBake to load `CCACHE_DIR` from the environment into the data store:

```
export BB_ENV_EXTRAWHITE="$BB_ENV_EXTRAWHITE CCACHE_DIR"
```

2. Tell BitBake to export what you have loaded into the environment store to the task environment of every running task. Loading something from the environment into the data store (previous step) only makes it available in the datastore. To export it to the task environment of every running task, use a command similar to the following in your `local.conf` or distro configuration file:

```
export CCACHE_DIR
```

Note

A side effect of the previous steps is that BitBake records the variable as a dependency of the build process in things like the shared state checksums. If doing so results in unnecessary rebuilds of tasks, you can whitelist the variable so that the shared state code ignores the dependency when it creates checksums. For information on this process, see the `BB_HASHBASE_WHITELIST` example in the "Checksums (Signatures)" section.

B.6. BitBake Command Line

Following is the BitBake help output:

```
$ bitbake --help
Usage: bitbake [options] [package ...]
```

Executes the specified task (default is 'build') for a given set of BitBake files. It expects that `BBFILES` is defined, which is a space separated list of files to be executed. `BBFILES` does support wildcards. Default `BBFILES` are the `.bb` files in the current directory.

Options:

<code>--version</code>	show program's version number and exit
<code>-h, --help</code>	show this help message and exit
<code>-b BUILDFILE, --buildfile=BUILDFILE</code>	execute the task against this <code>.bb</code> file, rather than a package from <code>BBFILES</code> . Does not handle any dependencies.
<code>-k, --continue</code>	continue as much as possible after an error. While the target that failed, and those that depend on it, cannot be remade, the other dependencies of these targets can be processed all the same.
<code>-a, --tryaltconfigs</code>	continue with builds by trying to use alternative providers where possible.
<code>-f, --force</code>	force run of specified cmd, regardless of stamp status

-c CMD, --cmd=CMD	Specify task to execute. Note that this only executes the specified task for the providee and the packages it depends on, i.e. 'compile' does not implicitly call stage for the dependencies (IOW: use only if you know what you are doing). Depending on the base.bbclass a listtasks tasks is defined and will show available tasks
-r PREFILE, --read=PREFILE	read the specified file before bitbake.conf
-R POSTFILE, --postread=POSTFILE	read the specified file after bitbake.conf
-v, --verbose	output more chit-chat to the terminal
-D, --debug	Increase the debug level. You can specify this more than once.
-n, --dry-run	don't execute, just go through the motions
-S, --dump-signatures	don't execute, just dump out the signature construction information
-p, --parse-only	quit after parsing the BB files (developers only)
-s, --show-versions	show current and preferred versions of all packages
-e, --environment	show the global or per-package environment (this is what used to be bbread)
-g, --graphviz	emit the dependency trees of the specified packages in the dot syntax
-I EXTRA_ASSUME_PROVIDED, --ignore-deps=EXTRA_ASSUME_PROVIDED	Assume these dependencies don't exist and are already provided (equivalent to ASSUME_PROVIDED). Useful to make dependency graphs more appealing
-l DEBUG_DOMAINS, --log-domains=DEBUG_DOMAINS	Show debug logging for the specified logging domains
-P, --profile	profile the command and print a report
-u UI, --ui=UI	userinterface to use
-t SERVERTYPE, --servertype=SERVERTYPE	Choose which server to use, none, process or xmlrpc
--revisions-changed	Set the exit code depending on whether upstream floating revisions have changed or not

B.7. Fetchers

BitBake also contains a set of "fetcher" modules that allow retrieval of source code from various types of sources. For example, BitBake can get source code from a disk with the metadata, from websites, from remote shell accounts or from Source Code Management (SCM) systems like cvs/subversion/git.

Fetchers are usually triggered by entries in SRC_URI. You can find information about the options and formats of entries for specific fetchers in the BitBake manual [<http://bitbake.berlios.de/manual/>].

One useful feature for certain Source Code Manager (SCM) fetchers is the ability to "auto-update" when the upstream SCM changes version. Since this ability requires certain functionality from the SCM, not all systems support it. Currently Subversion, Bazaar and to a limited extent, Git support the ability to "auto-update". This feature works using the SRCREV variable. See the "Development Within Yocto Project for a Package that Uses an External SCM" section for more information.

Appendix C. Reference: Classes

Class files are used to abstract common functionality and share it amongst multiple .bb files. Any metadata usually found in a .bb file can also be placed in a class file. Class files are identified by the extension .bbclass and are usually placed in a classes/ directory beneath the meta*/ directory found in the Yocto Project file's area. Class files can also be pointed to by BUILDDIR (e.g. build/) in the same way as .conf files in the conf directory. Class files are searched for in BBPATH using the same method by which .conf files are searched.

In most cases inheriting the class is enough to enable its features, although for some classes you might need to set variables or override some of the default behaviour.

C.1. The base class -**base.bbclass**

The base class is special in that every .bb file inherits it automatically. This class contains definitions for standard basic tasks such as fetching, unpacking, configuring (empty by default), compiling (runs any Makefile present), installing (empty by default) and packaging (empty by default). These classes are often overridden or extended by other classes such as autotools.bbclass or package.bbclass. The class also contains some commonly used functions such as oe_runmake.

C.2. Autotooled Packages -**autotools.bbclass**

Autotools (autoconf, automake, and libtool) bring standardization. This class defines a set of tasks (configure, compile etc.) that work for all Autotooled packages. It should usually be enough to define a few standard variables and then simply inherit autotools. This class can also work with software that emulates Autotools. For more information, see the "Autotooled Package [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#usingpoky-extend-addpkg-autotools>]" section in the Yocto Project Development Manual.

It's useful to have some idea of how the tasks defined by this class work and what they do behind the scenes.

- `do_configure` # regenerates the configure script (using autoreconf) and then launches it with a standard set of arguments used during cross-compilation. You can pass additional parameters to configure through the EXTRA_OECONF variable.
- `do_compile` # runs make with arguments that specify the compiler and linker. You can pass additional arguments through the EXTRA_OEMAKE variable.
- `do_install` # runs make install and passes a DESTDIR option, which takes its value from the standard DESTDIR variable.

C.3. Alternatives -**update-alternatives.bbclass**

Several programs can fulfill the same or similar function and be installed with the same name. For example, the ar command is available from the busybox, binutils and elfutils packages. The update-alternatives.bbclass class handles renaming the binaries so that multiple packages can be installed without conflicts. The ar command still works regardless of which packages are installed or subsequently removed. The class renames the conflicting binary in each package and symlinks the highest priority binary during installation or removal of packages.

Four variables control this class:

- ALTERNATIVE_NAME # The name of the binary that is replaced (ar in this example).
- ALTERNATIVE_LINK # The path to the resulting binary (/bin/ar in this example).
- ALTERNATIVE_PATH # The path to the real binary (/usr/bin/ar.binutils in this example).
- ALTERNATIVE_PRIORITY # The priority of the binary. The version with the most features should have the highest priority.

Currently, the Yocto Project supports only one binary per package.

C.4. Initscripts -**update-rc.d.bbclass**

This class uses `update-rc.d` to safely install an initialization script on behalf of the package. The Yocto Project takes care of details such as making sure the script is stopped before a package is removed and started when the package is installed. Three variables control this class: `INITSCRIPT_PACKAGES`, `INITSCRIPT_NAME` and `INITSCRIPT_PARAMS`. See the variable links for details.

C.5. Binary config scripts -**binconfig.bbclass**

Before `pkg-config` had become widespread, libraries shipped shell scripts to give information about the libraries and include paths needed to build software (usually named `LIBNAME-config`). This class assists any recipe using such scripts.

During staging, BitBake installs such scripts into the `sysroots/` directory. BitBake also changes all paths to point into the `sysroots/` directory so all builds that use the script will use the correct directories for the cross compiling layout.

C.6. Debian renaming -**debian.bbclass**

This class renames packages so that they follow the Debian naming policy (i.e. `eglibc` becomes `libc6` and `eglibc-devel` becomes `libc6-dev`).

C.7. Pkg-config -**pkgconfig.bbclass**

`pkg-config` brought standardization and this class aims to make its integration smooth for all libraries that make use of it.

During staging, BitBake installs `pkg-config` data into the `sysroots/` directory. By making use of `sysroot` functionality within `pkg-config`, this class no longer has to manipulate the files.

C.8. Distribution of sources - **src_distribute_local.bbclass**

Many software licenses require that source files be provided along with the binaries. To simplify this process, two classes were created: `src_distribute.bbclass` and `src_distribute_local.bbclass`.

The results of these classes are `tmp/deploy/source/` subdirs with sources sorted by `LICENSE` field. If recipes list few licenses (or have entries like "Bitstream Vera"), the source archive is placed in each license directory.

This class operates using three modes:

- `copy`: Copies the files to the distribute directory.
- `symlink`: Symlinks the files to the distribute directory.
- `move+symlink`: Moves the files into the distribute directory and then symlinks them back.

C.9. Perl modules -**cpan.bbclass**

Recipes for Perl modules are simple. These recipes usually only need to point to the source's archive and then inherit the proper `.bbclass` file. Building is split into two methods depending on which method the module authors used.

Modules that use old `Makefile.PL`-based build system require `cpan.bbclass` in their recipes.

Modules that use `Build.PL`-based build system require using `cpan_build.bbclass` in their recipes.

C.10. Python extensions -**distutils.bbclass**

Recipes for Python extensions are simple. These recipes usually only need to point to the source's archive and then inherit the proper `.bbclass` file. Building is split into two methods depending on which method the module authors used.

Extensions that use an Autotools-based build system require Autotools and `distutils-based.bbclass` files in their recipes.

Extensions that use `distutils`-based build systems require `distutils.bbclass` in their recipes.

C.11. Developer Shell -**devshell.bbclass**

This class adds the `devshell` task. Distribution policy dictates whether to include this class as the Yocto Project does. See the "Development Within a Development Shell" section for more information about using `devshell`.

C.12. Packaging -**package*.bbclass**

The packaging classes add support for generating packages from a build's output. The core generic functionality is in `package.bbclass`. The code specific to particular package types is contained in various sub-classes such as `package_deb.bbclass`, `package_ipk.bbclass`, and `package_rpm.bbclass`. Most users will want one or more of these classes.

You can control the list of resulting package formats by using the `PACKAGE_CLASSES` variable defined in the `local.conf` configuration file found in the Yocto Project file's `conf` directory. When defining the variable, you can specify one or more package types. Since images are generated from packages, a packaging class is needed to enable image generation. The first class listed in this variable is used for image generation.

The package class you choose can affect build-time performance and has space ramifications. In general, building a package with RPM takes about thirty percent more time as compared to using IPK to build the same or similar package. This comparison takes into account a complete build of the package with all dependencies previously built. The reason for this discrepancy is because the RPM package manager creates and processes more metadata than the IPK package manager. Consequently, you might consider setting `PACKAGE_CLASSES` to `"package_ipk"` if you are building smaller systems.

Keep in mind, however, that RPM starts to provide more abilities than IPK due to the fact that it processes more metadata. For example, this information includes individual file types, file checksum generation and evaluation on install, sparse file support, conflict detection and resolution for multilib systems, ACID style upgrade, and repackaging abilities for rollbacks.

Another consideration for packages built using the RPM package manager is space. For smaller systems, the extra space used for the Berkley Database and the amount of metadata can affect your ability to do on-device upgrades.

You can find additional information on the effects of the package class at these two Yocto Project mailing list links:

- <https://lists.yoctoproject.org/pipermail/poky/2011-May/006362.html> [<http://lists.yoctoproject.org/pipermail/poky/2011-May/006362.html>]
- <https://lists.yoctoproject.org/pipermail/poky/2011-May/006363.html> [<http://lists.yoctoproject.org/pipermail/poky/2011-May/006363.html>]

C.13. Building kernels -**kernel.bbclass**

This class handles building Linux kernels. The class contains code to build all kernel trees. All needed headers are staged into the `STAGING_KERNEL_DIR` directory to allow out-of-tree module builds using `module.bbclass`.

This means that each built kernel module is packaged separately and inter-module dependencies are created by parsing the `modinfo` output. If all modules are required, then installing the `kernel-`

modules package installs all packages with modules and various other kernel packages such as kernel-vmlinux.

Various other classes are used by the kernel and module classes internally including kernel-arch.bbclass, module_strip.bbclass, module-base.bbclass, and linux-kernel-base.bbclass.

C.14. Creating images -**image.bbclass** and **rootfs*.bbclass**

These classes add support for creating images in several formats. First, the root filesystem is created from packages using one of the rootfs_*.bbclass files (depending on the package format used) and then the image is created.

The IMAGE_FSTYPES variable controls the types of images to generate.

The IMAGE_INSTALL variable controls the list of packages to install into the image.

C.15. Host System sanity checks - **sanity.bbclass**

This class checks to see if prerequisite software is present so that users can be notified of potential problems that might affect their build. The class also performs basic user configuration checks from the local.conf configuration file to prevent common mistakes that cause build failures. Distribution policy usually whether to include this class as the Yocto Project does.

C.16. Generated output quality assurance checks - **insane.bbclass**

This class adds a step to the package generation process that sanity checks the packages generated by the Yocto Project. A range of checks are performed that check the build's output for common problems that show up during runtime. Distribution policy usually dictates whether to include this class as the Yocto Project does.

You can configure the sanity checks so that specific test failures either raise a warning or an error message. Typically, failures for new tests generate a warning. Subsequent failures for the same test would then generate an error message once the metadata is in a known and good condition. You use the WARN_QA variable to specify tests for which you want to generate a warning message on failure. You use the ERROR_QA variable to specify tests for which you want to generate an error message on failure.

The following list shows the tests you can list with the WARN_QA and ERROR_QA variables:

- **ldflags**: Ensures that the binaries were linked with the LDFLAGS options provided by the build system. If this test fails, check that the LDFLAGS variable is being passed to the linker command.
- **useless-rpaths**: Checks for dynamic library load paths (rpaths) in the binaries that by default on a standard system are searched by the linker (e.g. /lib and /usr/lib). While these paths will not cause any breakage, they do waste space and are unnecessary.
- **rpaths**: Checks for rpaths in the binaries that contain build system paths such as TMPDIR. If this test fails, bad -rpath options are being passed to the linker commands and your binaries have potential security issues.
- **dev-so**: Checks that the .so symbolic links are in the -dev package and not in any of the other packages. In general, these symlinks are only useful for development purposes. Thus, the -dev package is the correct location for them. Some very rare cases do exist for dynamically loaded modules where these symlinks are needed instead in the main package.
- **debug-files**: Checks for .debug directories in anything but the -dbg package. The debug files should all be in the -dbg package. Thus, anything packaged elsewhere is incorrect packaging.

- **arch**: Checks the Executable and Linkable Format (ELF) type, bit size and endianness of any binaries to ensure it matches the target architecture. This test fails if any binaries don't match the type since there would be an incompatibility. Sometimes software, like bootloaders, might need to bypass this check.
- **debug-deps**: Checks that `-dbg` packages only depend on other `-dbg` packages and not on any other types of packages, which would cause a packaging bug.
- **dev-deps**: Checks that `-dev` packages only depend on other `-dev` packages and not on any other types of packages, which would be a packaging bug.
- **pkgconfig**: Checks `.pc` files for any `TMPDIR`/`WORKDIR` paths. Any `.pc` file containing these paths is incorrect since `pkg-config` itself adds the correct `sysroot` prefix when the files are accessed.
- **la**: Checks `.la` files for any `TMPDIR` paths. Any `.la` file containing these paths is incorrect since `libtool` adds the correct `sysroot` prefix when using the files automatically itself.
- **desktop**: Runs the `desktop-file-validate` program against any `.desktop` files to validate their contents against the specification for `.desktop` files.

C.17. Autotools configuration data cache - **siteinfo.bbclass**

Autotools can require tests that must execute on the target hardware. Since this is not possible in general when cross compiling, site information is used to provide cached test results so these tests can be skipped over but still make the correct values available. The `meta/site` directory contains test results sorted into different categories such as architecture, endianness, and the `libc` used. Site information provides a list of files containing data relevant to the current build in the `CONFIG_SITE` variable that Autotools automatically picks up.

The class also provides variables like `SITEINFO_ENDIANNES` and `SITEINFO_BITS` that can be used elsewhere in the metadata.

Because this class is included from `base.bbclass`, it is always active.

C.18. Adding Users -**useradd.bbclass**

If you have packages that install files that are owned by custom users or groups, you can use this class to specify those packages and associate the users and groups with those packages. The `meta-skeleton/recipes-skeleton/useradd/useradd-example.bb` recipe in the Yocto Project Files provides a simple example that shows how to add three users and groups to two packages. See the `useradd-example.bb` for more information on how to use this class.

C.19. Using External Source - **externalsrc.bbclass**

You can use this class to build software from source code that is external to the Yocto Project build system. In other words, your source code resides in an external tree outside of the Yocto Project. Building software from an external source tree means that the normal fetch, unpack, and patch process is not used.

To use the class, you need to define the `S` variable to point to the directory that contains the source files. You also need to have your recipe inherit the `externalsrc.bbclass` class.

This class expects the source code to support recipe builds that use the `B` variable to point to the directory in which the Yocto Project build system places the generated objects built from the recipes. By default, the `B` directory is set to the following, which is separate from the source directory (`S`):

```
${WORKDIR}/${BPN}-${PV}/
```

See the glossary entries for the `WORKDIR`, `BPN`, `PV`, `S`, and `B` for more information.

You can build object files in the external tree by setting the B variable equal to "\${S}". However, this practice does not work well if you use the source for more than one variant (i.e., "natives" such as quilt-native, or "crosses" such as gcc-cross). So, be sure there are no "native", "cross", or "multilib" variants of the recipe.

If you do want to build different variants of a recipe, you can use the BBCLASSEXTEND variable. When you do, the B variable must support the recipe's ability to build variants in different working directories. Most autotools-based recipes support separating these directories. The Yocto Project defaults to using separate directories for gcc and some kernel recipes. Alternatively, you can make sure that separate recipes exist that each use the BBCLASSEXTEND variable to build each variant. The separate recipes can inherit a single target recipe.

For information on how to use this class, see the "Building Software from an External Source [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#building-software-from-an-external-source>]" section in the Yocto Project Development Manual.

C.20. Other Classes

Thus far, this appendix has discussed only the most useful and important classes. However, other classes exist within the meta/classes directory in the Yocto Project file's directory structure. You can examine the .bbclass files directly for more information.

Appendix D. Reference: Images

The Yocto Project build process supports several types of images to satisfy different needs. When you issue the bitbake command you provide a “top-level” recipe that essentially begins the build for the type of image you want.

Note

Building an image without GNU Public License Version 3 (GPLv3) components is only supported for minimal and base images. Furthermore, if you are going to build an image using non-GPLv3 components, you must make the following changes in the `local.conf` file before using the BitBake command to build the minimal or base image:

1. Comment out the `EXTRA_IMAGE_FEATURES` line
2. Set `INCOMPATIBLE_LICENSE = "GPLv3"`

From within the poky Git repository, use the following command to list the supported images:

```
$ ls meta*/recipes*/images/*.bb
```

These recipes reside in the `meta/recipes-core/images`, `meta/recipes-extended/images`, `meta/recipes-graphics/images`, and `meta/recipes-sato/images` directories of your local Yocto Project file structure (Git repository or extracted release tarball). Although the recipe names are somewhat explanatory, here is a list that describes them:

- `core-image-base`: A console-only image that fully supports the target device hardware.
- `core-image-core`: An X11 image with simple applications such as terminal, editor, and file manager.
- `core-image-minimal`: A small image just capable of allowing a device to boot.
- `core-image-minimal-dev`: A `core-image-minimal` image suitable for development work.
- `core-image-minimal-initramfs`: A `core-image-minimal` image that has the Minimal RAM-based Initial Root Filesystem (initramfs) as part of the kernel, which allows the system to find the first “init” program more efficiently.
- `core-image-minimal-mtdutils`: A `core-image-minimal` image that has support for the Minimal MTD Utilities, which let the user interact with the MTD subsystem in the kernel to perform operations on flash devices.
- `core-image-basic`: A foundational basic image without support for X that can be reasonably used for customization.
- `core-image-lsb`: A `core-image-basic` image suitable for implementations that conform to Linux Standard Base (LSB).
- `core-image-lsb-dev`: A `core-image-lsb` image that is suitable for development work.
- `core-image-lsb-sdk`: A `core-image-lsb` that includes everything in `meta-toolchain` but also includes development headers and libraries to form a complete standalone SDK. See the “External Development Using the Meta-Toolchain” section for more information.
- `core-image-clutter`: An image with support for the Open GL-based toolkit Clutter, which enables development of rich and animated graphical user interfaces.
- `core-image-sato`: An image with Sato support, a mobile environment and visual style that works well with mobile devices. The image supports X11 with a Sato theme and Pimlico applications and also contains terminal, editor, and file manager.
- `core-image-sato-dev`: A `core-image-sato` image suitable for development that also includes a native toolchain and libraries needed to build applications on the device itself. The image also

includes testing and profiling tools as well as debug symbols. This image was formerly `core-image-sdk`.

- `core-image-sato-sdk`: A `core-image-sato` image that includes everything in meta-toolchain. The image also includes development headers and libraries to form a complete standalone SDK. See the "External Development Using the Meta-Toolchain" section for more information.
- `core-image-rt`: A `core-image-minimal` image plus a real-time test suite and tools appropriate for real-time use.
- `core-image-rt-sdk`: A `core-image-rt` image that includes everything in meta-toolchain. The image also includes development headers and libraries to form a complete stand-alone SDK. See the "External Development Using the Meta-Toolchain" section for more information.
- `core-image-gtk-directfb`: An image that uses gtk+ over directfb instead of X11. In order to build, this image requires specific distro configuration that enables gtk over directfb.
- `self-hosted-image`: An image you can run using the Yocto Project Build Appliance. The image is suitable to load and boot from either VMware Player [<http://www.vmware.com/products/player/overview.html>] or VMWare Workstation [<http://www.vmware.com/products/workstation/overview.html>]. For more information on the Build Appliance, see the Build Appliance [<http://www.yoctoproject.org/documentation/build-appliance>] page on the Yocto Project website.

Tip

From the Yocto Project release 1.1 onwards, `-live` and `-directdisk` images have been replaced by a "live" option in `IMAGE_FSTYPES` that will work with any image to produce an image file that can be copied directly to a CD or USB device and run as is. To build a live image, simply add "live" to `IMAGE_FSTYPES` within the `local.conf` file or wherever appropriate and then build the desired image as normal.

Appendix E. Reference: Features

Features provide a mechanism for working out which packages should be included in the generated images. Distributions can select which features they want to support through the `DISTRO_FEATURES` variable, which is set in the `poky.conf` distribution configuration file. Machine features are set in the `MACHINE_FEATURES` variable, which is set in the machine configuration file and specifies the hardware features for a given machine.

These two variables combine to work out which kernel modules, utilities, and other packages to include. A given distribution can support a selected subset of features so some machine features might not be included if the distribution itself does not support them.

E.1. Distro

The items below are valid options for `DISTRO_FEATURES`:

- `alsa`: ALSA support will be included (OSS compatibility kernel modules will be installed if available).
- `bluetooth`: Include bluetooth support (integrated BT only)
- `ext2`: Include tools for supporting for devices with internal HDD/Microdrive for storing files (instead of Flash only devices)
- `irda`: Include Irda support
- `keyboard`: Include keyboard support (e.g. keymaps will be loaded during boot).
- `pci`: Include PCI bus support
- `pcmcia`: Include PCMCIA/CompactFlash support
- `usb gadget`: USB Gadget Device support (for USB networking/serial/storage)
- `usb host`: USB Host support (allows to connect external keyboard, mouse, storage, network etc)
- `wifi`: WiFi support (integrated only)
- `cramfs`: CramFS support
- `ipsec`: IPSec support
- `ipv6`: IPv6 support
- `nfs`: NFS client support (for mounting NFS exports on device)
- `ppp`: PPP dialup support
- `smbfs`: SMB networks client support (for mounting Samba/Microsoft Windows shares on device)

E.2. Machine

The items below are valid options for `MACHINE_FEATURES`:

- `acpi`: Hardware has ACPI (x86/x86_64 only)
- `alsa`: Hardware has ALSA audio drivers
- `apm`: Hardware uses APM (or APM emulation)
- `bluetooth`: Hardware has integrated BT
- `ext2`: Hardware HDD or Microdrive
- `irda`: Hardware has Irda support
- `keyboard`: Hardware has a keyboard

- pci: Hardware has a PCI bus
- pcmcia: Hardware has PCMCIA or CompactFlash sockets
- screen: Hardware has a screen
- serial: Hardware has serial support (usually RS232)
- touchscreen: Hardware has a touchscreen
- usb gadget: Hardware is USB gadget device capable
- usb host: Hardware is USB Host capable
- wifi: Hardware has integrated WiFi

E.3. Reference: Images

The contents of images generated by the Yocto Project can be controlled by the `IMAGE_FEATURES` and `EXTRA_IMAGE_FEATURES` variables that you typically configure in your image recipes. Through these variables you can add several different predefined packages such as development utilities or packages with debug information needed to investigate application problems or profile applications.

Current list of `IMAGE_FEATURES` contains the following:

- apps-console-core: Core console applications such as `ssh`, `daemon`, `avahi daemon`, `portmap` (for mounting NFS shares)
- x11-base: X11 server + minimal desktop
- x11-sato: OpenedHand Sato environment
- apps-x11-core: Core X11 applications such as an X Terminal, file manager, and file editor
- apps-x11-games: A set of X11 games
- apps-x11-pimlico: OpenedHand Pimlico application suite
- tools-sdk: A full SDK that runs on the device
- tools-debug: Debugging tools such as `strace` and `gdb`
- tools-profile: Profiling tools such as `oprofile`, `exmap`, and `LTTng`
- tools-testapps: Device testing tools (e.g. touchscreen debugging)
- nfs-server: NFS server (exports / over NFS to everybody)
- dev-pkgs: Development packages (headers and extra library links) for all packages installed in a given image
- dbg-pkgs: Debug packages for all packages installed in a given image

Appendix F. Reference: Variables Glossary

This section lists common variables used in the Yocto Project and gives an overview of their function and contents.

Glossary

A B C D E F H I K L M P R S T W

A

ALLOW_EMPTY

Specifies if an output package should still be produced if it is empty. By default, BitBake does not produce empty packages. This default behavior can cause issues when there is an RDEPENDS or some other runtime hard-requirement on the existence of the package.

Like all package-controlling variables, you must always use them in conjunction with a package name override. Here is an example:

```
ALLOW_EMPTY_${PN}
```

AUTHOR

The email address used to contact the original author or authors in order to send patches, forward bugs, etc.

AUTOREV

Specifies to use the current (newest) source revision. This variable is with the SRCREV variable.

B

B

The directory in which the Yocto Project build system places generated objects during a recipe's build process. By default, this directory is the same as the S directory:

```
B = ${WORKDIR}/${BPN}-${PV}/
```

You can separate the source directory (S) and the directory pointed to by the B variable. Most autotools-based recipes support separating these directories. The Yocto Project defaults to using separate directories for gcc and some kernel recipes.

BAD_RECOMMENDATIONS

A list of packages not to install despite being recommended by a recipe. Support for this variable exists only for images that use the ipkg packaging system.

BBCLASSEXTEND

Allows you to extend a recipe so that it builds variants of the software. Common variants for recipes exist such as "natives" like quilt-native, which is a copy of quilt built to run on the build system; "crosses" such as gcc-cross, which is a compiler built to run on the build machine but produces binaries that run on the target MACHINE; "nativesdk", which targets the SDK machine instead of MACHINE; and "multilibs" in the form "multilib:<multilib_name>".

To build a different variant of the recipe with a minimal amount of code, it usually is as simple as adding the following to your recipe:

```
BBCLASSEXTEND += "native nativesdk"  
BBCLASSEXTEND += "multilib:<multilib_name>"
```

BBMASK	Prevents BitBake from processing recipes and recipe append files. You can use the BBMASK variable to "hide" these .bb and .bbappend files. BitBake ignores any recipe or recipe append files that match the expression. It is as if BitBake does not see them at all. Consequently, matching files are not parsed or otherwise used by BitBake. The value you provide is passed to python's regular expression compiler. For complete syntax information, see python's documentation at http://docs.python.org/release/2.3/lib/re-syntax.html. The expression is compared against the full paths to the files. For example, the following uses a complete regular expression to tell BitBake to ignore all recipe and recipe append files in the <code>./meta-ti/recipes-misc/</code> directory: <pre>BBMASK = ".*meta-ti/recipes-misc/"</pre> Use the BBMASK variable from within the <code>conf/local.conf</code> file found in the Yocto Project build directory.
BB_NUMBER_THREADS	The maximum number of tasks BitBake should run in parallel at any one time. If your host development system supports multiple cores a good rule of thumb is to set this variable to twice the number of cores.
BBFILE_COLLECTIONS	Lists the names of configured layers. These names are used to find the other BBFILE_* variables. Typically, each layer will append its name to this variable in its <code>conf/layer.conf</code> file.
BBFILE_PATTERN	Variable that expands to match files from BBFILES in a particular layer. This variable is used in the <code>conf/layer.conf</code> file and must be suffixed with the name of the specific layer (e.g. BBFILE_PATTERN_emenlow).
BBFILE_PRIORITY	Assigns the priority for recipe files in each layer. This variable is useful in situations where the same package appears in more than one layer. Setting this variable allows you to prioritize a layer against other layers that contain the same package - effectively letting you control the precedence for the multiple layers. The precedence established through this variable stands regardless of a layer's package version (PV variable). For example, a layer that has a package with a higher PV value but for which the BBFILE_PRIORITY is set to have a lower precedence still has a lower precedence. A larger value for the BBFILE_PRIORITY variable results in a higher precedence. For example, the value 6 has a higher precedence than the value 5. If not specified, the BBFILE_PRIORITY variable is set based on layer dependencies (see the LAYERDEPENDS variable for more information. The default priority, if unspecified for a layer with no dependencies, is the lowest defined priority + 1 (or 1 if no priorities are defined). Tip You can use the command <code>bitbake-layers show_layers</code> to list all configured layers along with their priorities.
BBFILES	List of recipe files used by BitBake to build software
BBPATH	Used by BitBake to locate .bbclass and configuration files. This variable is analogous to the PATH variable.
BBINCLUDELOGS	Variable that controls how BitBake displays logs on build failure.

BBLAYERS Lists the layers to enable during the Yocto Project build. This variable is defined in the `bblayers.conf` configuration file in the Yocto Project build directory. Here is an example:

```
BBLAYERS = " \
/home/scottrif/poky/meta \
/home/scottrif/poky/meta-yocto \
/home/scottrif/poky/meta-mykernel \
"
```

This example enables three layers, one of which is a custom, user-defined layer named `meta-mykernel`.

BPN Bare name of package with any suffixes like `-cross` `-native` removed.

C

CFLAGS Flags passed to C compiler for the target system. This variable evaluates to the same as `TARGET_CFLAGS`.

COMPATIBLE_MACHINE A regular expression which evaluates to match the machines the recipe works with. It stops recipes being run on machines for which they are not compatible. This is particularly useful with kernels. It also helps to increase parsing speed as further parsing of the recipe is skipped if it is found the current machine is not compatible.

CONFFILES Identifies editable or configurable files that are part of a package. If the Package Management System (PMS) is being used to update packages on the target system, it is possible that configuration files you have changed after the original installation and that you now want to remain unchanged are overwritten. In other words, editable files might exist in the package that you do not want reset as part of the package update process. You can use the `CONFFILES` variable to list the files in the package that you wish to prevent the PMS from overwriting during this update process.

To use the `CONFFILES` variable, provide a package name override that identifies the package. Then, provide a space-separated list of files. Here is an example:

```
CONFFILES_${PN} += "${sysconfdir}/file1 \
${sysconfdir}/file2 ${sysconfdir}/file3"
```

A relationship exists between the `CONFFILES` and `FILES` variables. The files listed within `CONFFILES` must be a subset of the files listed within `FILES`. Because the configuration files you provide with `CONFFILES` are simply being identified so that the PMS will not overwrite them, it makes sense that the files must already be included as part of the package through the `FILES` variable.

Note

When specifying paths as part of the `CONFFILES` variable, it is good practice to use appropriate path variables. For example, `${sysconfdir}` rather than `/etc` or `${bindir}` rather than `/usr/bin`. You can find a list of these variables at the top of the `/meta/conf/bitbake.conf` file in the Yocto Project files directory.

CONFIG_SITE A list of files that contains `autoconf` test results relevant to the current build. This variable is used by the Autotools utilities when running `configure`.

CORE_IMAGE_EXTRA_INSTALL Specifies the list of packages to be added to the image. This variable should only be set in the `local.conf` configuration file found in the Yocto Project Build Directory [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#yocto-project-build-directory>].

This variable replaces `POKY_EXTRA_INSTALL`, which is no longer supported.

D

D The destination directory.

DEBUG_BUILD Specifies to build packages with debugging information. This influences the value of the `SELECTED_OPTIMIZATION` variable.

DEBUG_OPTIMIZATION The options to pass in `TARGET_CFLAGS` and `CFLAGS` when compiling a system for debugging. This variable defaults to `"-O -fno-omit-frame-pointer -g"`.

DEFAULT_PREFERENCE Specifies the priority of recipes.

DEPENDS A list of build-time dependencies for a given recipe. The variable indicates recipes that must have been staged before a particular recipe can configure.

DESCRIPTION The package description used by package managers.

DESTDIR the destination directory.

DISTRO The short name of the distribution.

DISTRO_EXTRA_RRECOMMENDS The list of packages which extend usability of the image. Those packages will automatically be installed but can be removed by user.

DISTRO_FEATURES The features of the distribution.

DISTRO_NAME The long name of the distribution.

DISTRO_PN_ALIAS Alias names used for the recipe in various Linux distributions.

See the "Handling a Package Name Alias" [http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#usingpoky-configuring-DISTRO_PN_ALIAS] section in the Yocto Project Development Manual for more information.

DISTRO_VERSION the version of the distribution.

DL_DIR The central download directory used by the build process to store downloads. You can set this directory by defining the `DL_DIR` variable in the `/conf/local.conf` file. This directory is self-maintaining and you should not have to touch it. By default, the directory is `downloads` in the Yocto Project build directory.

```
#DL_DIR ?= "${TOPDIR}/downloads"
```

To specify a different download directory, simply uncomment the line and provide your directory.

During a first build, the system downloads many different source code tarballs from various upstream projects. Downloading can take a while, particularly if your network connection is slow. Tarballs are all stored in the directory defined by `DL_DIR` and the build system looks there first to find source tarballs.

Note

When wiping and rebuilding, you can preserve this directory to speed up this part of subsequent builds.

You can safely share this directory between multiple builds on the same development machine. For additional information on how the build process gets source files when working behind a firewall or proxy server, see the "FAQ" appendix.

E

ENABLE_BINARY_LOCALE_GENERATION

Variable that controls which locales for eglibc are to be generated during the build (useful if the target device has 64Mbytes of RAM or less).

EXTRA_IMAGE_FEATURES

Allows extra packages to be added to the generated images. You set this variable in the `local.conf` configuration file. Note that some image features are also added using the `IMAGE_FEATURES` variable generally configured in image recipes. You can use this variable to add more features in addition to those. Here are some examples of features you can add:

"dbg-pkgs" - Adds -dbg packages for all installed packages including symbol information for debugging and profiling.

"dev-pkgs" - Adds -dev packages for all installed packages. This is useful if you want to develop against the libraries in the image.

"tools-sdk" - Adds development tools such as gcc, make, pkgconfig and so forth.

"tools-debug" - Adds debugging tools such as gdb and strace.

"tools-profile" - Adds profiling tools such as oprofile, exmap, lttng and valgrind (x86 only).

"tools-testapps" - Adds useful testing tools such as ts_print, aplay, arecord and so forth.

"debug-tweaks" - Makes an image suitable for development. For example, ssh root access has a blank password. You should remove this feature before you produce a production image.

There are other application targets too, see `meta/classes/poky-image.bbclass` and `meta/packages/tasks/task-poky.bb` for more details.

EXTRA_IMAGEDEPENDS

A list of recipes to be built that do not provide packages to be installed in the root filesystem.

Sometimes a recipe is required to build the final image but is not needed in the root filesystem. You can use the `EXTRA_IMAGEDEPENDS` variable to list these recipes and thus, specify the dependencies. A typical example is a required bootloader in a machine configuration.

Note

To add packages to the root filesystem, see the various `*DEPENDS` and `*RECOMMENDS` variables.

`EXTRA_OECMAKE`

Additional cmake options.

`EXTRA_OECONF`

Additional configure script options.

`EXTRA_OEMAKE`

Additional GNU make options.

F

`FILES`

The list of directories or files that are placed in packages.

To use the `FILES` variable, provide a package name override that identifies the package. Then, provide a space-separated list of files or paths that identifies the files you want included as part of the package. Here is an example:

```
FILES_${PN} += "${bindir}/mydir1/ ${bindir}/mydir2/myfile"
```

Note

When specifying paths as part of the `FILES` variable, it is good practice to use appropriate path variables. For example, `${sysconfdir}` rather than `/etc` or `${bindir}` rather than `/usr/bin`. You can find a list of these variables at the top of the `/meta/conf/bitbake.conf` file in the Yocto Project files directory.

If some of the files you provide with the `FILES` variable are editable and you know they should not be overwritten during the package update process by the Package Management System (PMS), you can identify these files so that the PMS will not overwrite them. See the `CONFFILES` variable for information on how to identify these files to the PMS.

`FILESEXTRAPATHS`

Extends the search path the Yocto Project build system uses when looking for files and patches as it processes recipes. The directories BitBake uses when it processes recipes is defined by the `FILESPATH` variable. You can add directories to the search path by defining the `FILESEXTRAPATHS` variable.

To add paths to the search order, provide a list of directories and separate each path using a colon character as follows:

```
FILESEXTRAPATHS_prepend := "path_1:path_2:path_3:"
```

Typically, you want your directories searched first. To make sure that happens, use `_prepend` and the immediate expansion (`:=`) operator as shown in the previous example. Finally, to maintain the integrity of the `FILESPATH` variable, you must include the appropriate beginning or ending (as needed) colon character.

The `FILESEXTRAPATHS` variable is intended for use in `.bbappend` files to include any additional files provided in that layer. You typically accomplish this with the following:

```
FILESEXTRAPATHS_prepend := "${THISDIR}/${PN}:"
```

FILESPATH	The default set of directories the Yocto Project build system uses when searching for patches and files. During the build process, BitBake searches each directory in FILESPATH in the specified order when looking for files and patches specified by each file:// URI in a recipe. The default value for the FILESPATH variable is defined in the <code>base.bbclass</code> class found in <code>meta/classes</code> in the Yocto Project Files [http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#yocto-project-files]: <pre>FILESPATH = "\${@base_set_filespath(["\${FILE_DIRNAME}/\${PF}", \ "\${FILE_DIRNAME}/\${P}", "\${FILE_DIRNAME}/\${PN}", \ "\${FILE_DIRNAME}/\${BP}", "\${FILE_DIRNAME}/\${BPN}", \ "\${FILE_DIRNAME}/files", "\${FILE_DIRNAME}"], d)}</pre> Do not hand-edit the FILESPATH variable. If you want to extend the set of pathnames that BitBake uses when searching for files and patches, use the FILESETRAPATHS variable.
FILESYSTEM_PERMS_TABLES	Allows you to define your own file permissions settings table as part of your configuration for the packaging process. For example, suppose you need a consistent set of custom permissions for a set of groups and users across an entire work project. It is best to do this in the packages themselves but this is not always possible. By default, the Yocto Project uses the <code>fs-perms.txt</code>, which is located in the <code>meta/files</code> directory of the Yocto Project files directory. If you create your own file permissions setting table, you should place it in your layer or the distros layer. You define the FILESYSTEM_PERMS_TABLES variable in the <code>conf/local.conf</code> file, which is found in the Yocto Project's build directory, to point to your custom <code>fs-perms.txt</code>. You can specify more than a single file permissions setting table. The paths you specify to these files must be defined within the BBPATH variable. For guidance on how to create your own file permissions settings table file, examine the existing <code>fs-perms.txt</code>.
FULL_OPTIMIZATION	The options to pass in TARGET_CFLAGS and CFLAGS when compiling an optimized system. This variable defaults to <code>"-fexpensive-optimizations -fomit-frame-pointer -frename-registers -O2"</code>.
H	
Homepage	Website where more info about package can be found
I	
IMAGE_FEATURES	The list of features present in images. Typically, you configure this variable in image recipes. Note that you can add extra features to the image by using the EXTRA_IMAGE_FEATURES variable. See the "Reference: Images" section for the list of features present in images built by the Yocto Project.
IMAGE_FSTYPES	Formats of root filesystem images that you want to have created.
IMAGE_INSTALL	Specifies the packages to install into an image. The IMAGE_INSTALL variable is a mechanism for an image recipe and you should use it with care to avoid ordering issues.

Image recipes set `IMAGE_INSTALL` to specify the packages to install into an image through `image.bbclass`. Additionally, "helper" classes exist, such as `core-image.bbclass`, that can take `IMAGE_FEATURES` lists and turn these into auto-generated entries in `IMAGE_INSTALL` in addition to its default contents.

Using `IMAGE_INSTALL` with the `+=` operator from the `/conf/local.conf` file or from within an image recipe is not recommended as it can cause ordering issues. Since `core-image.bbclass` sets `IMAGE_INSTALL` to a default value using the `?=` operator, using a `+=` operation against `IMAGE_INSTALL` will result in unexpected behavior when used in `/conf/local.conf`. Furthermore, the same operation from within an image recipe may or may not succeed depending on the specific situation. In both these cases, the behavior is contrary to how most users expect the `+=` operator to work.

When you use this variable, it is best to use it as follows:

```
IMAGE_INSTALL_append = " package-name"
```

Be sure to include the space between the quotation character and the start of the package name.

`IMAGE_OVERHEAD_FACTOR`

Defines a multiplier that the build system applies to the initial image size for cases when the multiplier times the returned disk usage value for the image is greater than the sum of `IMAGE_ROOTFS_SIZE` and `IMAGE_ROOTFS_EXTRA_SPACE`. The result of the multiplier applied to the initial image size creates free disk space in the image as overhead. By default, the build process uses a multiplier of 1.3 for this variable. This default value results in 30% free disk space added to the image when this method is used to determine the final generated image size. You should be aware that post install scripts and the package management system uses disk space inside this overhead area. Consequently, the multiplier does not produce an image with all the theoretical free disk space. See `IMAGE_ROOTFS_SIZE` for information on how the build system determines the overall image size.

The default 30% free disk space typically gives the image enough room to boot and allows for basic post installs while still leaving a small amount of free disk space. If 30% free space is inadequate, you can increase the default value. For example, the following setting gives you 50% free space added to the image:

```
IMAGE_OVERHEAD_FACTOR = "1.5"
```

Alternatively, you can ensure a specific amount of free disk space is added to the image by using `IMAGE_ROOTFS_EXTRA_SPACE` the variable.

`IMAGE_ROOTFS_EXTRA_SPACE`

Defines additional free disk space created in the image in Kbytes. By default, this variable is set to "0". This free disk space is added to the image after the build system determines the image size as described in `IMAGE_ROOTFS_SIZE`.

This variable is particularly useful when you want to ensure that a specific amount of free disk space is available on a device after an image is installed and running. For example, to be sure 5 Gbytes of free disk space is available, set the variable as follows:

IMAGE_ROOTFS_EXTRA_SPACE = "5242880"

IMAGE_ROOTFS_SIZE

Defines the size in Kbytes for the generated image. The Yocto Project build system determines the final size for the generated image using an algorithm that takes into account the initial disk space used for the generated image, a requested size for the image, and requested additional free disk space to be added to the image. Programatically, the build system determines the final size of the generated image as follows:

```
if (image-du * overhead) < rootfs-size:
internal-rootfs-size = rootfs-size + xspace
else:
internal-rootfs-size = (image-du * overhead) + xspace
```

where:

image-du = Returned value of the du command on the image.

overhead = IMAGE_OVERHEAD_FACTOR

rootfs-size = IMAGE_ROOTFS_SIZE

internal-rootfs-size = Initial root filesystem size before any modifications.

xspace = IMAGE_ROOTFS_EXTRA_SPACE

INC_PR

Defines the Package revision. You manually combine values for INC_PR into the PR field of the parent recipe. When you change this variable, you change the PR value for every person that includes the file.

The following example shows how to use the INC_PR variable given a common .inc file that defines the variable. Once defined, you can use the variable to set the PR value:

```
recipes-graphics/xorg-font/font-util_1.1.1.bb:PR - "${INC_PR}.1"
recipes-graphics/xorg-font/xorg-font-common.inc:INC_PR - "r1"
recipes-graphics/xorg-font/encondings_1.0.3.bb:PR - "${INC_PR}.1"
recipes-graphics/xorg-font/fiont-alias_1.0.2.bb:PR - "${INC_PR}.0"
```

INHIBIT_PACKAGE_STRIP

Causes the build to not strip binaries in resulting packages.

INHERIT

Causes the named class to be inherited at this point during parsing. The variable is only valid in configuration files.

INITSCRIPT_PACKAGES

A list of the packages that contain initscripts. If multiple packages are specified, you need to append the package name to the other INITSCRIPT_* as an override.

This variable is used in recipes when using update-rc.d.bbclass. The variable is optional and defaults to the PN variable.

INITSCRIPT_NAME

The filename of the initscript (as installed to \${etcdir}/init.d).

This variable is used in recipes when using update-rc.d.bbclass. The variable is Mandatory.

INITSCRIPT_PARAMS

Specifies the options to pass to update-rc.d. An example is start 99 5 2 . stop 20 0 1 6 ., which gives the script a runlevel of 99,

starts the script in initlevels 2 and 5, and stops the script in levels 0, 1 and 6.

The variable is mandatory and is used in recipes when using update-rc.d.bbclass.

K

KERNEL_FEATURES

Includes additional metadata from the Linux Yocto kernel Git repository. In the Yocto Project build system, the default Board Support Packages (BSPs) metadata is provided through the KMACHINE and KBRANCH variables. You can use the KERNEL_FEATURES variable to further add metadata for all BSPs.

The metadata you add through this variable includes config fragments and features descriptions, which usually includes patches as well as config fragments. You typically override the KERNEL_FEATURES variable for a specific machine. In this way, you can provide validated, but optional, sets of kernel configurations and features.

For example, the following adds netfilter to all the Linux Yocto kernels and adds sound support to the qemu86 machine:

```
# Add netfilter to all linux-yocto kernels
KERNEL_FEATURES="features/netfilter"

# Add sound support to the qemu86 machine
KERNEL_FEATURES_append_qemu86="cfg/sound"
```

KERNEL_IMAGETYPE

The type of kernel to build for a device, usually set by the machine configuration files and defaults to "zImage". This variable is used when building the kernel and is passed to make as the target to build.

L

LAYERDEPENDS

Lists the layers that this recipe depends upon, separated by spaces. Optionally, you can specify a specific layer version for a dependency by adding it to the end of the layer name with a colon, (e.g. "anotherlayer:3" to be compared against LAYERVERSION_anotherlayer in this case). An error will be produced if any dependency is missing or the version numbers do not match exactly (if specified). This variable is used in the conf/layer.conf file and must be suffixed with the name of the specific layer (e.g. LAYERDEPENDS_mylayer).

LAYERDIR

When used inside the layer.conf configuration file, this variable provides the path of the current layer. This variable requires immediate expansion (see the BitBake manual) as lazy expansion can result in the expansion happening in the wrong directory and therefore giving the wrong value.

LAYERVERSION

Optionally specifies the version of a layer as a single number. You can use this within LAYERDEPENDS for another layer in order to depend on a specific version of the layer. This variable is used in the conf/layer.conf file and must be suffixed with the name of the specific layer (e.g. LAYERVERSION_mylayer).

LIC_FILES_CHKSUM

Checksums of the license text in the recipe source code.

This variable tracks changes in license text of the source code files. If the license text is changed, it will trigger a build failure, which gives the developer an opportunity to review any license change.

This variable must be defined for all recipes (unless LICENSE is set to "CLOSED")

For more information, see the [Tracking License Changes](#) section

LICENSE

The list of package source licenses.

LICENSE_DIR

Path to additional licenses used during the build. By default, the Yocto Project uses COMMON_LICENSE_DIR to define the directory that holds common license text used during the build. The LICENSE_DIR variable allows you to extend that location to other areas that have additional licenses:

```
LICENSE_DIR += "/path/to/additional/common/licenses"
```

M

MACHINE

Specifies the target device.

MACHINE_ESSENTIAL_EXTRA_RDEPENDS

A list of required packages to install as part of the package being built. The build process depends on these packages being present. Furthermore, because this is a "machine essential" variable, the list of packages are essential for the machine to boot. The impact of this variable affects images based on task-core-boot, including the core-image-minimal image.

This variable is similar to the MACHINE_ESSENTIAL_EXTRA_RRECOMMENDS variable with the exception that the package being built has a build dependency on the variable's list of packages. In other words, the image will not build if a file in this list is not found.

For example, suppose you are building a runtime package that depends on a certain disk driver. In this case, you would use the following:

```
MACHINE_ESSENTIAL_EXTRA_RDEPENDS += "<disk_driver>"
```

MACHINE_ESSENTIAL_EXTRA_RRECOMMENDS

A list of recommended packages to install as part of the package being built. The build process does not depend on these packages being present. Furthermore, because this is a "machine essential" variable, the list of packages are essential for the machine to boot. The impact of this variable affects images based on task-core-boot, including the core-image-minimal image.

This variable is similar to the MACHINE_ESSENTIAL_EXTRA_RDEPENDS variable with the exception that the package being built does not have a build dependency on the variable's list of packages. In other words, the image will build if a file in this list is not found. However, because this is one of the "essential" variables, the resulting image might not boot on the machine. Or, if the machine does boot using the image, the machine might not be fully functional.

Consider an example where you have a custom kernel with a disk driver built into the kernel itself, rather than using the driver built as a module. If you include the package that has the driver module as part of the variable's list, the build process will not find that package. However, because these packages are "recommends" packages, the

build will not fail due to the missing package. Not accounting for any other problems, the custom kernel would still boot the machine.

Some example packages of these machine essentials are flash, screen, keyboard, mouse, or touchscreen drivers (depending on the machine).

For example, suppose you are building a runtime package that depends on a mouse driver. In this case, you would use the following:

```
MACHINE_ESSENTIAL_EXTRA_RRECOMMENDS += "<mouse_driver>"
```

MACHINE_EXTRA_RDEPENDS A list of optional but non-machine essential packages to install as part of the package being built. Even though these packages are not essential for the machine to boot, the build process depends on them being present. The impact of this variable affects all images based on task-base, which does not include the core-image-minimal or core-image-basic images.

This variable is similar to the MACHINE_EXTRA_RRECOMMENDS variable with the exception that the package being built has a build dependency on the variable's list of packages. In other words, the image will not build if a file in this list is not found.

An example is a machine that might or might not have a WiFi card. The package containing the WiFi support is not essential for the machine to boot the image. If it is not there, the machine will boot but not be able to use the WiFi functionality. However, if you include the package with the WiFi support as part of the variable's package list, the build process depends on finding the package. In this case, you would use the following:

```
MACHINE_EXTRA_RDEPENDS += "<wifi_driver>"
```

MACHINE_EXTRA_RRECOMMENDS A list of optional but non-machine essential packages to install as part of the package being built. The package being built has no build dependency on the list of packages with this variable. The impact of this variable affects only images based on task-base, which does not include the core-image-minimal or core-image-basic images.

This variable is similar to the MACHINE_EXTRA_RDEPENDS variable with the exception that the package being built does not have a build dependency on the variable's list of packages. In other words, the image will build if a file in this list is not found.

An example is a machine that might or might not have a WiFi card. The package containing the WiFi support is not essential for the machine to boot the image. If it is not there, the machine will boot but not be able to use the WiFi functionality. You are free to either include or not include the the package with the WiFi support as part of the variable's package list, the build process does not depend on finding the package. If you include the package, you would use the following:

```
MACHINE_EXTRA_RRECOMMENDS += "<wifi_driver>"
```

MACHINE_FEATURES Specifies the list of device features. See the Machine section for more information.

MAINTAINER The email address of the distribution maintainer.

P

PACKAGE_ARCH	The architecture of the resulting package.
PACKAGE_CLASSES	This variable, which is set in the <code>local.conf</code> configuration file found in the Yocto Project file's <code>conf</code> directory, specifies the package manager to use when packaging data. You can provide one or more arguments for the variable with the first argument being the package manager used to create images: <pre>PACKAGE_CLASSES ?= "package_rpm package_deb package_ipk"</pre> For information on build performance effects as a result of the package manager use, see <code>Packaging - package*.bbclass</code> in this manual.
PACKAGE_EXTRA_ARCHS	Specifies the list of architectures compatible with the device CPU. This variable is useful when you build for several different devices that use miscellaneous processors such as XScale and ARM926-EJS).
PACKAGES	The list of packages to be created from the recipe. The default value is <code>"\${PN}-dbg \${PN} \${PN}-doc \${PN}-dev"</code> .
PARALLEL_MAKE	Specifies extra options that are passed to the <code>make</code> command during the compile tasks. This variable is usually in the form <code>-j 4</code> , where the number represents the maximum number of parallel threads <code>make</code> can run. If your development host supports multiple cores a good rule of thumb is to set this variable to twice the number of cores on the host.
PN	The name of the package.
PR	The revision of the package. The default value for this variable is <code>"r0"</code> .
PV	The version of the package. The version is normally extracted from the recipe name. For example, if the recipe is named <code>expat_2.0.1.bb</code> , then <code>PV</code> will be <code>2.0.1</code> . <code>PV</code> is generally not overridden within a recipe unless it is building an unstable version from a source code repository (e.g. Git or Subversion).
PE	the epoch of the package. The default value is <code>"0"</code> . The field is used to make upgrades possible when the versioning scheme changes in some backwards incompatible way.
PREFERRED_PROVIDER	If multiple recipes provide an item, this variable determines which recipe should be given preference. The variable must always be suffixed with the name of the provided item, and should be set to the <code>\$PN</code> of the recipe to which you want to give precedence. Here is an example: <pre>PREFERRED_PROVIDER_virtual/xserver = "xserver-xf86"</pre>
PREFERRED_VERSION	If there are multiple versions of recipes available, this variable determines which recipe should be given preference. The variable must always be suffixed with the <code>\$PN</code> for which to select, and should be set to the <code>\$PV</code> to which you want to give precedence. You can use the <code>"%"</code> character as a wildcard to match any number of characters, which can be useful when specifying versions that contain long revision numbers that could potentially change. Here are two examples:

```
PREFERRED_VERSION_python = "2.6.6"  
PREFERRED_VERSION_linux-yocto = "3.0+git%"
```

R

RCONFLICTS The list of packages that conflict with this package. Note that the package will not be installed if the conflicting packages are not first removed.

RDEPENDS A list of packages that must be installed as part of a package being built. The package being built has a runtime dependency on the packages in the variable's list. In other words, in order for the package being built to run correctly, it depends on these listed packages. If a package in this list cannot be found during the build, the build will not complete.

Because the RDEPENDS variable applies to packages being built, you should always attach an override to the variable to specify the particular runtime package that has the dependency. For example, suppose you are building a development package that depends on the perl package. In this case, you would use the following RDEPENDS statement:

```
RDEPENDS_${PN}-dev += "perl"
```

In the example, the package name (\${PN}-dev) must appear as it would in the PACKAGES namespace before any renaming of the output package by classes like `debian.bbclass`.

Some automatic handling occurs around the RDEPENDS variable:

- **shlibdeps:** If a runtime package contains a shared library (.so), the build processes the library in order to determine other libraries to which it is dynamically linked. The build process adds these libraries to RDEPENDS to create the runtime package.
- **pcdeps:** If the package ships a pkg-config information file, the build process uses this file to add items to the RDEPENDS variable to create the runtime packages.

RRECOMMENDS A list of packages that extend the usability of a package being built. The package being built does not depend on this list of packages in order to successfully build, but needs them for the extended usability. To specify runtime dependencies for packages, see the RDEPENDS variable.

The Yocto Project build process automatically installs the list of packages as part of the built package. However, you can remove them later if you want. If, during the build, a package from the list cannot be found, the build process continues without an error.

Because the RRECOMMENDS variable applies to packages being built, you should always attach an override to the variable to specify the particular package whose usability is being extended. For example, suppose you are building a development package that is extended to support wireless functionality. In this case, you would use the following:

```
RRECOMMENDS_${PN}-dev += "<wireless_package_name>"
```

In the example, the package name (`${PN}-dev`) must appear as it would in the `PACKAGES` namespace before any renaming of the output package by classes like `debian.bbclass`.

RREPLACES

The list of packages that are replaced with this package.

S

S

The location in the Yocto Project Build Directory [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#yocto-project-build-directory>] where unpacked package source code resides. This location is within the working directory (`WORKDIR`), which is not static. The unpacked source location depends on the package name (`PN`) and package version (`PV`) as follows:

```
${WORKDIR}/${PN}-${PV}
```

As an example, assume a Yocto Project Files [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#yocto-project-files>] top-level directory named `poky` and a default Yocto Project Build Directory of `poky/build`. In this case, the working directory the build system uses to build the `db` package is the following:

```
~/poky/build/tmp/work/qemux86-poky-linux/db-5.1.19-r3/db-5.1.19
```

SECTION

The section where package should be put. Package managers use this variable.

SELECTED_OPTIMIZATION

The variable takes the value of `FULL_OPTIMIZATION` unless `DEBUG_BUILD = "1"`. In this case the value of `DEBUG_OPTIMIZATION` is used.

SERIAL_CONSOLE

The speed and device for the serial port used to attach the serial console. This variable is given to the kernel as the `"console"` parameter and after booting occurs `getty` is started on that port so remote login is possible.

SSTATE_DIR

The directory for the shared state.

SITEINFO_ENDIANNESS

Specifies the endian byte order of the target system. The variable is either `"le"` for little-endian or `"be"` for big-endian.

SITEINFO_BITS

Specifies the number of bits for the target system CPU. The variable is either `"32"` or `"64"`.

SRC_URI

The list of source files - local or remote.

SRC_URI_OVERRIDES_PACKAGE_ARCH

By default, the Yocto Project automatically detects whether `SRC_URI` contains files that are machine-specific. If so, the Yocto Project automatically changes `PACKAGE_ARCH`. Setting this variable to `"0"` disables this behavior.

SRCDATE

The date of the source code used to build the package. This variable applies only if the source was fetched from a Source Code Manager (SCM).

SRCREV

The revision of the source code used to build the package. This variable applies to Subversion, Git, Mercurial and Bazaar only. Note

that if you wish to build a fixed revision and you wish to avoid performing a query on the remote repository every time BitBake parses your recipe, you should specify a SRCREV that is a full revision identifier and not just a tag.

STAGING_KERNEL_DIR	The directory with kernel headers that are required to build out-of-tree modules.
STAMP	The directory (usually TMPDIR/stamps) with timestamps of executed tasks.
SUMMARY	The short (72 characters or less) summary of the binary package for packaging systems such as ipkg, rpm or debian. By default, this variable inherits DESCRIPTION.

T

TARGET_ARCH	The architecture of the device being built. While a number of values are possible, the Yocto Project primarily supports arm and i586.
TARGET_CFLAGS	Flags passed to the C compiler for the target system. This variable evaluates to the same as CFLAGS.
TARGET_FPU	Specifies the method for handling FPU code. For FPU-less targets, which include most ARM CPUs, the variable must be set to "soft". If not, the kernel emulation gets used, which results in a performance penalty.
TARGET_OS	Specifies the target's operating system. The variable can be set to "linux" for glibc-based systems and to "linux-uclibc" for uclibc. For ARM/EABI targets, there are also "linux-gnueabi" and "linux-uclibc-gnueabi" values possible.
TCLIBC	Specifies which variant of the GNU standard C library (libc) to use during the build process. This variable replaces POKYLIBC, which is no longer supported. You can select eglibc or uclibc.

Note

This release of the Yocto Project does not support the glibc implementation of libc.

TCMODE	The toolchain selector. This variable replaces POKYMODE, which is no longer supported. The TCMODE variable selects the external toolchain built from the Yocto Project or a few supported combinations of the upstream GCC or CodeSourcery Labs toolchain. The variable determines which of the files in meta/conf/distro/include/tcmode-* is used. By default, TCMODE is set to "default", which chooses tcmode-default.inc. The variable is similar to TCLIBC, which controls the variant of the GNU standard C library (libc) used during the build process: eglibc or uclibc.
TMPDIR	This variable is the temporary directory the Yocto Project build system uses when it does its work building images. By default, the TMPDIR variable is named tmp within the Yocto Project Build Directory [http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#yocto-project-build-directory]. If you want to establish this directory in a location other than the default, you can uncomment the following statement in the conf/local.conf file in the Yocto

Project Files [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#yocto-project-files>]:

```
#TMPDIR = "${TOPDIR}/tmp"
```

TOPDIR

This variable is the Yocto Project Build Directory [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#yocto-project-build-directory>]. BitBake automatically sets this variable. The Yocto Project build system uses the build directory when building images.

W

WORKDIR

The pathname of the working directory in which the Yocto Project build system builds packages. This directory is located within the TMPDIR directory structure and changes as different packages are built.

The actual WORKDIR directory depends on several things:

- The temporary directory - TMPDIR
- The package architecture - PACKAGE_ARCH
- The target machine - MACHINE
- The target operating system - TARGET_OS
- The package name - PN
- The package version - PV
- The package revision - PR

For packages that are not dependent on a particular machine, WORKDIR is defined as follows:

```
${TMPDIR}/work/${PACKAGE_ARCH}-poky-${TARGET_OS}/${PN}-${PV}-${PR}
```

As an example, assume a Yocto Project Files [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#yocto-project-files>] top-level directory named poky and a default Yocto Project Build Directory [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#yocto-project-build-directory>] of poky/build. In this case, the working directory the build system uses to build the v86d package is the following:

```
~/poky/build/tmp/work/qemux86-poky-linux/v86d-01.9-r0
```

For packages that are dependent on a particular machine, WORKDIR is defined slightly different:

```
${TMPDIR}/work/${MACHINE}-poky-${TARGET_OS}/${PN}-${PV}-${PR}
```

As an example, again assume a Yocto Project Files top-level directory named poky and a default Yocto Project build directory of poky/build. In this case, the working directory the build system uses to

build the `acl` package, which is dependent on a MIPS-based device, is the following:

```
~/poky/build/tmp/work/mips-poky-linux/acl-2.2.51-r2
```

Appendix G. Reference: Variable Context

While most variables can be used in almost any context such as `.conf`, `.bbclass`, `.inc`, and `.bb` files, some variables are often associated with a particular locality or context. This appendix describes some common associations.

G.1. Configuration

The following subsections provide lists of variables whose context is configuration: distribution, machine, and local.

G.1.1. Distribution (Distro)

This section lists variables whose context is the distribution, or distro.

- DISTRO
- DISTRO_NAME
- DISTRO_VERSION
- MAINTAINER
- PACKAGE_CLASSES
- TARGET_OS
- TARGET_FPU
- TCMODE
- TCLIBC

G.1.2. Machine

This section lists variables whose context is the machine.

- TARGET_ARCH
- SERIAL_CONSOLE
- PACKAGE_EXTRA_ARCHS
- IMAGE_FSTYPES
- MACHINE_FEATURES
- MACHINE_EXTRA_RDEPENDS
- MACHINE_EXTRA_RRECOMMENDS
- MACHINE_ESSENTIAL_EXTRA_RDEPENDS
- MACHINE_ESSENTIAL_EXTRA_RRECOMMENDS

G.1.3. Local

This section lists variables whose context is the local configuration through the `local.conf` file.

- DISTRO
- MACHINE

- DL_DIR
- BBFILES
- EXTRA_IMAGE_FEATURES
- PACKAGE_CLASSES
- BB_NUMBER_THREADS
- BBINCLUDELOGS
- ENABLE_BINARY_LOCALE_GENERATION

G.2. Recipes

The following subsections provide lists of variables whose context is recipes: required, dependencies, path, and extra build information.

G.2.1. Required

This section lists variables that are required for recipes.

- DESCRIPTION
- LICENSE
- LIC_FILES_CHKSUM
- SECTION
- HOMEPAGE
- AUTHOR
- SRC_URI

G.2.2. Dependencies

This section lists variables that define recipe dependencies.

- DEPENDS
- RDEPENDS
- RRECOMMENDS
- RCONFLICTS
- RREPLACES

G.2.3. Paths

This section lists variables that define recipe paths.

- WORKDIR
- S
- FILES

G.2.4. Extra Build Information

This section lists variables that define extra build information for recipes.

- DISTRO_PN_ALIAS

- EXTRA_OECMAKE
- EXTRA_OECONF
- EXTRA_OEMAKE
- PACKAGES
- DEFAULT_PREFERENCE

Appendix H. FAQ

H.1. How does Poky differ from OpenEmbedded [<http://www.openembedded.org>]?

Poky is the Yocto Project build system that was derived from OpenEmbedded [<http://www.openembedded.org>]. Poky is a stable, smaller subset focused on the mobile environment. Development in the Yocto Project using Poky is closely tied to OpenEmbedded with features being merged regularly between the two for mutual benefit.

H.2. I only have Python 2.4 or 2.5 but BitBake requires Python 2.6 or 2.7. Can I still use the Yocto Project?

You can use a stand-alone tarball to provide Python 2.6. You can find pre-built 32 and 64-bit versions of Python 2.6 at the following locations:

- 32-bit tarball [<http://downloads.yoctoproject.org/releases/miscsupport/python-nativesdk-standalone-i686.tar.bz2>]
- 64-bit tarball [http://downloads.yoctoproject.org/releases/miscsupport/python-nativesdk-standalone-x86_64.tar.bz2]

These tarballs are self-contained with all required libraries and should work on most Linux systems. To use the tarballs extract them into the root directory and run the appropriate command:

```
$ export PATH=/opt/poky/sysroots/i586-pokysdk-linux/usr/bin:$PATH
$ export PATH=/opt/poky/sysroots/x86_64-pokysdk-linux/usr/bin:$PATH
```

Once you run the command, BitBake uses Python 2.6.

H.3. How can you claim Poky is stable?

There are three areas that help with stability;

- The Yocto Project team keeps Poky small and focused. It contains around 650 packages as compared to over 5000 for full OpenEmbedded.
- The Yocto Project only supports hardware that the team has access to for testing.
- The Yocto Project uses an autobuilder, which provides continuous build and integration tests.

H.4. How do I get support for my board added to the Yocto Project?

There are two main ways to get a board supported in the Yocto Project;

- Send the Yocto Project team information on the board and if the team does not have it yet they will consider adding it.
- Send the Yocto Project team the BitBake recipes if you have them.

Usually, if the board is not completely exotic, adding support in the Yocto Project is fairly straightforward.

H.5. Are there any products using Poky?

The Vernier LabQuest [<http://vernier.com/labquest/>] is using the Yocto Project build system Poky. See the Vernier LabQuest [<http://www.vernier.com/products/interfaces/labq/>] for more information. There are a number of pre-production devices using Poky and the Yocto Project team announces them as soon as they are released.

H.6. What does the Yocto Project build system Poky produce as output?

Because the same set of recipes can be used to create output of various formats, the output of a Yocto Project build depends on how it was started. Usually, the output is a flashable image ready for the target device.

H.7. How do I add my package to the Yocto Project?

To add a package, you need to create a BitBake recipe. For information on how to add a package, see the section "Adding a Package [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#usingpoky-extend-addpkg>]" in the Yocto Project Development Manual.

H.8. Do I have to reflash my entire board with a new Yocto Project image when recompiling a package?

The Yocto Project can build packages in various formats such as ipk for ipkg/opkg, Debian package (.deb), or RPM. The packages can then be upgraded using the package tools on the device, much like on a desktop distribution such as Ubuntu or Fedora.

H.9. What is GNOME Mobile and what is the difference between GNOME Mobile and GNOME?

GNOME Mobile is a subset of the GNOME [<http://www.gnome.org>] platform targeted at mobile and embedded devices. The the main difference between GNOME Mobile and standard GNOME is that desktop-orientated libraries have been removed, along with deprecated libraries, creating a much smaller footprint.

H.10. I see the error 'chmod: XXXXX new permissions are r-xrwxrwx, not r-xr-xr-x'. What is wrong?

You are probably running the build on an NTFS filesystem. Use ext2, ext3, or ext4 instead.

H.11. How do I make the Yocto Project work in RHEL/CentOS?

To get the Yocto Project working under RHEL/CentOS 5.1 you need to first install some required packages. The standard CentOS packages needed are:

- "Development tools" (selected during installation)
- texi2html
- compat-gcc-34

On top of these, you need the following external packages:

- python-sqlite2 from DAG repository [<http://dag.wieers.com/rpm/packages/python-sqlite2/>]
- help2man from Karan repository [http://centos.karan.org/el4/extras/stable/x86_64/RPMS/repodata/repoview/help2man-0-1.33.1-2.html]

Once these packages are installed, the Yocto Project will be able to build standard images. However, there might be a problem with the QEMU emulator segfaulting. You can either disable the generation of binary locales by setting `ENABLE_BINARY_LOCALE_GENERATION` to "0" or by removing the `linux-2.6-execshield.patch` from the kernel and rebuilding it since that is the patch that causes the problems with QEMU.

H.12. I see lots of 404 responses for files on <http://www.yoctoproject.org/sources/> *. Is something wrong?

Nothing is wrong. The Yocto Project checks any configured source mirrors before downloading from the upstream sources. The Yocto Project does this searching for both source archives and pre-checked out versions of SCM managed software. These checks help in large installations because it can reduce load on the SCM servers themselves. The address above is one of the default mirrors configured into the Yocto Project. Consequently, if an upstream source disappears, the team can place sources there so builds continue to work.

H.13. I have machine-specific data in a package for one machine only but the package is being marked as machine-specific in all cases, how do I prevent this?

Set `SRC_URI_OVERRIDES_PACKAGE_ARCH = "0"` in the .bb file but make sure the package is manually marked as machine-specific in the case that needs it. The code that handles `SRC_URI_OVERRIDES_PACKAGE_ARCH` is in `base.bbclass`.

H.14. I'm behind a firewall and need to use a proxy server. How do I do that?

Most source fetching by the Yocto Project is done by `wget` and you therefore need to specify the proxy settings in a `.wgetrc` file in your home directory. Example settings in that file would be

```
http_proxy = http://proxy.yoyodyne.com:18023/  
ftp_proxy = http://proxy.yoyodyne.com:18023/
```

The Yocto Project also includes a `site.conf.sample` file that shows how to configure CVS and Git proxy servers if needed.

H.15. I'm using Ubuntu Intrepid and am seeing build failures. What's wrong?

In Intrepid, Ubuntu turns on by default the normally optional compile-time security features and warnings. There are more details at <https://wiki.ubuntu.com/CompilerFlags>. You can work around this problem by disabling those options by adding the following to the `BUILD_CPPFLAGS` variable in the `conf/bitbake.conf` file.

```
" -Wno-format-security -U_FORTIFY_SOURCE"
```

H.16. What's the difference between `foo` and `foo-native`?

The `*-native` targets are designed to run on the system being used for the build. These are usually tools that are needed to assist the build in some way such as `quilt-native`, which is used to apply patches. The non-native version is the one that runs on the target device.

H.17. I'm seeing random build failures. Help?!

If the same build is failing in totally different and random ways, the most likely explanation is that either the hardware you're running the build on has some problem, or, if you are running the build under virtualisation, the virtualisation probably has bugs. The Yocto Project processes a massive amount of data causing lots of network, disk and CPU activity and is sensitive to even single bit failures in any of these areas. True random failures have always been traced back to hardware or virtualisation issues.

H.18. What do we need to ship for license compliance?

This is a difficult question and you need to consult your lawyer for the answer for your specific case. It is worth bearing in mind that for GPL compliance there needs to be enough information shipped to allow someone else to rebuild the same end result you are shipping. This means sharing the source code, any patches applied to it, and also any configuration information about how that package was configured and built.

H.19. How do I disable the cursor on my touchscreen device?

You need to create a form factor file as described in the "Miscellaneous Recipe Files" section and set the `HAVE_TOUCHSCREEN` variable equal to one as follows:

```
HAVE_TOUCHSCREEN=1
```

H.20. How do I make sure connected network interfaces are brought up by default?

The default interfaces file provided by the `netbase` recipe does not automatically bring up network interfaces. Therefore, you will need to add a BSP-specific `netbase` that includes an interfaces file. See the "Miscellaneous Recipe Files" section for information on creating these types of miscellaneous recipe files.

For example, add the following files to your layer:

```
meta-MACHINE/recipes-bsp/netbase/netbase/MACHINE/interfaces  
meta-MACHINE/recipes-bsp/netbase/netbase_4.44.bbappend
```

H.21. How do I create images with more free space?

Images are created to be 1.2 times the size of the populated root filesystem. To modify this ratio so that there is more free space available, you need to set the configuration value `IMAGE_OVERHEAD_FACTOR`. For example, setting `IMAGE_OVERHEAD_FACTOR` to 1.5 sets the image size ratio to one and a half times the size of the populated root filesystem.

```
IMAGE_OVERHEAD_FACTOR = "1.5"
```

H.22. Why don't you support directories with spaces in the pathnames?

The Yocto Project team has tried to do this before but too many of the tools the Yocto Project depends on such as `autoconf` break when they find spaces in pathnames. Until that situation changes, the team will not support spaces in pathnames.

H.23. How do I use an external toolchain?

The toolchain configuration is very flexible and customizable. It is primarily controlled with the `TCMODE` variable. This variable controls which file to include (`conf/distro/include/tcmode-*.inc`).

The default value of `TCMODE` is "default". However, other patterns are accepted. In particular, "external-*" refers to external toolchains of which there are some basic examples included with the core. A user can use their own custom toolchain definition in their own layer (or as defined in the `local.conf` file) at the location `conf/distro/include/tcmode-*.inc`.

In addition to the toolchain configuration, you also need a corresponding toolchain recipe file. This recipe file needs to package up any pre-built objects in the toolchain such as `libgcc`, `libstdc++`, any locales and `libc`. An example is the `external-csl-toolchain_2008q3-72.bb`, which reuses the core `libc` packaging class to do most of the work.

H.24. How does the Yocto Project build system obtain source code and will it work behind my firewall or proxy server?

The way the Yocto Project obtains source code is highly configurable. You can setup the Yocto Project to get source code in most environments if HTTP transport is available.

When the build system searches for source code, it first tries the local download directory. If that location fails, Poky tries `PREMIRRORS`, the upstream source, and then `MIRRORS` in that order.

By default, Poky uses the Yocto Project source `PREMIRRORS` for SCM-based sources, upstreams for normal tarballs, and then falls back to a number of other mirrors including the Yocto Project source mirror if those fail.

As an example, you could add a specific server for Poky to attempt before any others by adding something like the following to the `local.conf` configuration file:

```
PREMIRRORS_prepend = "\n\ngit://.*/*.* http://www.yoctoproject.org/sources/ \n \n\nftp://.*/*.* http://www.yoctoproject.org/sources/ \n \n\nhttp://.*/*.* http://www.yoctoproject.org/sources/ \n \n\nhttps://.*/*.* http://www.yoctoproject.org/sources/ \n"
```

These changes cause Poky to intercept Git, FTP, HTTP, and HTTPS requests and direct them to the `http://` sources mirror. You can use `file://` URLs to point to local directories or network shares as well.

Aside from the previous technique, these options also exist:

```
BB_NO_NETWORK = "1"
```

This statement tells BitBake to throw an error instead of trying to access the Internet. This technique is useful if you want to ensure code builds only from local sources.

Here is another technique:

```
BB_FETCH_PREMIRRORONLY = "1"
```

This statement limits Poky to pulling source from the PREMIRRORS only. Again, this technique is useful for reproducing builds.

Here is another technique:

```
BB_GENERATE_MIRROR_TARBALLS = "1"
```

This statement tells Poky to generate mirror tarballs. This technique is useful if you want to create a mirror server. If not, however, the technique can simply waste time during the build.

Finally, consider an example where you are behind an HTTP-only firewall. You could make the following changes to the `local.conf` configuration file as long as the PREMIRROR server is up to date:

```
PREMIRRORS_prepend = "\n\
ftp://.*/* http://www.yoctoproject.org/sources/ \n \
http://.*/* http://www.yoctoproject.org/sources/ \n \
https://.*/* http://www.yoctoproject.org/sources/ \n"
BB_FETCH_PREMIRRORONLY = "1"
```

These changes would cause Poky to successfully fetch source over HTTP and any network accesses to anything other than the PREMIRROR would fail.

Poky also honors the standard shell environment variables `http_proxy`, `ftp_proxy`, `https_proxy`, and `all_proxy` to redirect requests through proxy servers.

Appendix I. Contributing to the Yocto Project

I.1. Introduction

The Yocto Project team is happy for people to experiment with the Yocto Project. A number of places exist to find help if you run into difficulties or find bugs. To find out how to download source code, see the "Yocto Project Release [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#local-yp-release>]" list item in the Yocto Project Development Manual.

I.2. Tracking Bugs

If you find problems with the Yocto Project, you should report them using the Bugzilla application at <http://bugzilla.yoctoproject.org>.

I.3. Mailing lists

To subscribe to the Yocto Project mailing lists, click on the following URLs and follow the instructions:

- <http://lists.yoctoproject.org/listinfo/yocto-announce>: Use this list to receive official Yocto Project announcements for developments and to learn about Yocto Project milestones.
- <http://lists.yoctoproject.org/listinfo/yocto>: Use this list to monitor Yocto Project development discussions, ask questions, and get help.
- <http://lists.yoctoproject.org/listinfo/poky>: Use this list to monitor discussions about the Yocto Project build system Poky, ask questions, and get help.

I.4. Internet Relay Chat (IRC)

Two IRC channels on freenode are available for the Yocto Project and Poky discussions:

- #yocto
- #poky

I.5. Links

Following is a list of resources you will find helpful:

- The Yocto Project website [<http://www.yoctoproject.org>]: The home site for the Yocto Project.
- OpenedHand [<http://o-hand.com>]: The company where the Yocto Project build system Poky was first developed. OpenedHand has since been acquired by Intel Corporation.
- Intel Corporation [<http://www.intel.com/>]: The company who acquired OpenedHand in 2008 and continues development on the Yocto Project.
- OpenEmbedded [<http://www.openembedded.org>]: The upstream, generic, embedded distribution the Yocto Project build system (Poky) derives from and to which it contributes.
- BitBake [<http://developer.berlios.de/projects/bitbake/>]: The tool used to process Yocto Project metadata.
- BitBake User Manual [<http://bitbake.berlios.de/manual/>]: A comprehensive guide to the BitBake tool.
- Pimlico [<http://pimlico-project.org/>]: A suite of lightweight Personal Information Management (PIM) applications designed primarily for handheld and mobile devices.

- QEMU [<http://wiki.qemu.org/Index.html>]: An open source machine emulator and virtualizer.

I.6. Contributions

The Yocto Project gladly accepts contributions. You can submit changes to the project either by creating and sending pull requests, or by submitting patches through email. For information on how to do both, see the "How to Submit a Change [<http://www.yoctoproject.org/docs/1.2.2/dev-manual/dev-manual.html#how-to-submit-a-change>]" section in the Yocto Project Development Manual.