
Fuego Documentation

Release 1.5.9

Cogent Embedded, Tim Bird

Apr 28, 2023

CONTENTS:

1	About Fuego	3
2	Architecture	5
2.1	Major elements	6
2.2	Container	6
2.3	Jenkins	6
2.4	Pre-packaged tests	7
2.5	Fuego test definition	7
2.6	Board	8
2.7	Fuego core	8
2.8	Different objects in Fuego	9
2.9	Jenkins operations	9
2.10	Test execution	10
2.11	Test environment file generation	10
2.12	Fuego test phases	11
2.13	Command line tool	15
3	Install and First Test	17
3.1	Overview	17
3.2	Step details	17
3.3	Run a test	19
3.4	Conclusions	19
3.5	What do do next?	20
4	Fuego Quickstart Guide	21
4.1	Overview	21
4.2	Install pre-requisite software	21
4.3	Download, build, start and access	22
4.4	Add your board to fuego	23
4.5	Install a toolchain	24
4.6	Now select some tests	24
4.7	Run a test	25
4.8	Additional Notes	25
4.9	Troubleshooting	26
5	Installing Fuego	27
5.1	Overview	27
5.2	Install pre-requisite software	27
5.3	Overview of remaining steps	28
5.4	Install the Fuego repositories	28

5.5	Create the Fuego container	29
5.6	Start the Fuego container	31
5.7	Access the Fuego Jenkins web interface	31
5.8	Access the Fuego docker command line	32
5.9	Remaining steps	33
5.10	Alternative installation configurations	33
6	Adding a Board	35
6.1	Overview	35
6.2	Board-specific test variables	39
7	Adding a toolchain	41
7.1	Introduction	41
7.2	Obtain a toolchain	41
7.3	Install the SDK in the docker container	42
7.4	Create a -tools.sh file for the toolchain	43
7.5	Reference the toolchain in a board file	44
7.6	Notes	44
8	Adding test jobs to Jenkins	45
8.1	Selecting tests or plans	45
8.2	Adding individual tests	45
8.3	Adding jobs based on testplans	46
9	Adding views to Jenkins	49
9.1	Adding a board view	49
9.2	Add view by test name regular expression	50
9.3	Add specific jobs	50
10	Test variables	51
10.1	Introduction	51
10.2	Board file	51
10.3	Overlay system	52
10.4	Stored variables	53
10.5	Spec variables	55
10.6	Dynamic variables	55
10.7	Variable precedence	55
11	Jenkins User interface	57
11.1	Main dashboard	57
11.2	Node pages	60
11.3	Job pages	60
11.4	Build pages	63
11.5	View pages	65
11.6	Other Jenkins pages	66
12	Command Line Tool - ftc	69
12.1	Introduction	69
12.2	Commands groups	70
12.3	ftc run-test	77
13	Log files	81
13.1	Main results logs	81
13.2	Other test artifacts	82
13.3	Summary	84

14	Generating Reports	85
14.1	ftc gen-report	85
14.2	ftc gen-report usage help	94
15	Adding or Customizing a Distribution	95
15.1	Introduction	95
15.2	Distribution overlay file	95
15.3	Referencing the distribution in the board file	96
15.4	Testing Fuego/distribution interactions	96
15.5	Notes	97
16	Building Documentation	99
16.1	building the outdated PDF	99
16.2	building the RST docs	100
17	Fuego Developer Notes	101
17.1	Resources	101
17.2	Notes	102
17.3	Logs	109
17.4	Core scripts	110
18	License And Contribution Policy	113
18.1	License	113
18.2	Submitting contributions	115
19	Integration with ttc	117
19.1	Outline of supported functionality	117
19.2	Supported operations	118
19.3	Location of ttc.conf	118
19.4	Steps to use ttc with a target board	118
19.5	modify your copy_to_cmd	119
20	Working with remote boards	121
20.1	using a jump server	121
20.2	Using ttc transport remotely	121
20.3	Setting up ssh ProxyCommand in the Fuego docker container	122
21	API Reference	123
22	Functions available for tests to call	125
22.1	Functions for interacting with the target	125
22.2	Functions for checking dependencies and requirements	125
22.3	Functions for preparing board and executing tests	126
22.4	Functions for parsing results	126
22.5	Functions for cleanup	126
22.6	For batch tests	126
22.7	Misceleneous functions	126
22.8	For printing messages at various message output levels	126
22.9	fuego_test.sh functions	127
23	Parser module API	129
23.1	Parser API	129
23.2	Deprecated API	130
23.3	Developer notes	131

24	Core interfaces	135
24.1	From Jenkins to Fuego	135
24.2	From Fuego to Fuego	137
24.3	Example Values	138
24.4	From Fuego to Jenkins	139
25	Adding a new test	143
25.1	Overview of Steps	143
25.2	Decide on a test name	143
25.3	Create the directory for the test	144
25.4	Get the source for a test	144
25.5	Test script	145
25.6	Test spec	146
25.7	Test results parser	147
25.8	Pass criteria and reference info	147
25.9	Jenkins job definition file	148
25.10	Publishing the test	148
25.11	Technical Details	148
26	Using Batch Tests	151
26.1	How to make a batch test	151
26.2	Test output	153
26.3	Preparing the system for a batch job	153
26.4	Executing a batch test	154
26.5	Viewing batch test results	154
26.6	Miscellaneous notes	154
27	Fuego naming rules	157
27.1	Fuego test name	157
27.2	Test files	157
27.3	Test spec names	158
27.4	Board names	158
27.5	Jenkins element names	158
27.6	Run identifier	158
27.7	timestamp	159
27.8	test identifiers	159
27.9	Test variable names	160
28	Artwork	163
28.1	Logos	163
28.2	Banners	164
28.3	images	164
28.4	Photos	165
28.5	Diagrams	165
28.6	Presentation templates	166
29	FAQ	167
29.1	Languages and formats used	167
30	Glossary	169
30.1	B	169
30.2	C	170
30.3	D	170
30.4	F	170
30.5	J	170

30.6	L	171
30.7	M	171
30.8	O	171
30.9	P	171
30.10	R	172
30.11	S	172
30.12	T	172
30.13	V	173
31	Raspberry Pi Fuego Setup	175
31.1	Obtain your network address	175
31.2	Configure the SSH server	176
31.3	Make a test directory	178
31.4	Add the board file to Fuego	179
31.5	Add the toolchain to Fuego	179
31.6	Add a node and jobs for the board	179
31.7	Run a board check	180
32	Using the qemuarm target	181
32.1	Build a qemuarm image	181
32.2	Running the qemuarm image	182
32.3	Test connectivity	182
32.4	Test building software	182
33	Variables	183
33.1	A	183
33.2	B	183
33.3	C	184
33.4	F	184
33.5	G	185
33.6	I	185
33.7	L	186
33.8	M	186
33.9	N	187
33.10	O	187
33.11	P	187
33.12	R	187
33.13	S	188
33.14	T	188
33.15	U	189
33.16	UNDOCUMENTED (YET)	189
34	Dynamic Variables	191
34.1	Notes	191
35	Troubleshooting Guide	193
35.1	Installation	193
35.2	General	195
35.3	Handling different Fuego Error messages	195
36	OSS Test Vision	197
36.1	Overview of concepts	197
36.2	Ideas related to the vision	202
36.3	Test definition or attributes	202
36.4	Test app store	203

36.5	Authenticating tests	203
36.6	Test system metrics	204
37	Test Specs and Plans	205
37.1	Introduction	205
37.2	Test plans	206
37.3	Test Specs	207
37.4	Variable use during test execution	208
37.5	Specifying failure cases	208
37.6	Catalog of current plans	209
38	parser.py	211
38.1	PROGRAM	211
38.2	DESCRIPTION	211
38.3	SAMPLES	212
38.4	ENVIRONMENT and ARGUMENTS	212
38.5	SOURCE	213
38.6	SEE ALSO	213
39	criteria.json	215
39.1	Introduction	215
39.2	Evaluation criteria	216
39.3	Customizing the criteria.json file for a board	217
39.4	Examples	218
39.5	Schema	219
39.6	Compatibility with previous Fuego versions	221
40	tools.sh	225
40.1	Variable usage details	226
41	Coding Style	227
41.1	Indentation and line length	227
41.2	Trailing whitespace	227
41.3	Shell features	228
41.4	Python style	228
42	Testplan Reference	229
42.1	Filename and location	229
42.2	Top level attributes	229
42.3	Individual test definitions	230
42.4	Test setting precedence	232
43	Indices and tables	233
43.1	Sandbox	233
43.2	Page Level Header (H1)	233
43.3	Sndbx2	237
43.4	Page Level Header2 (H1)	237
43.5	cmd	239
43.6	run.json	240
43.7	Fuego Test System	246
43.8	FUEGO BUILD FLAGS	248
43.9	FUEGO DEBUG	248
43.10	FUEGO LOGLEVELS	249

Note: This documentation is being updated with material from the Fuego wiki (at fuegotest.org). Please be patient while this work is in progress.

#.. include:: FrontPage.rst

ABOUT FUEGO

Fuego is a test system specifically designed for embedded Linux testing. It supports automated testing of embedded targets from a host system, as it's primary method of test execution.

The quick introduction to Fuego is that it consists of a host/target script engine and over 100 pre-packaged tests. These are installed in a docker container along with a Jenkins web interface and job control system, ready for out-of-the-box Continuous Integration testing of your embedded Linux project.

The idea is that in the simplest case, you just add your board, a toolchain, and go!

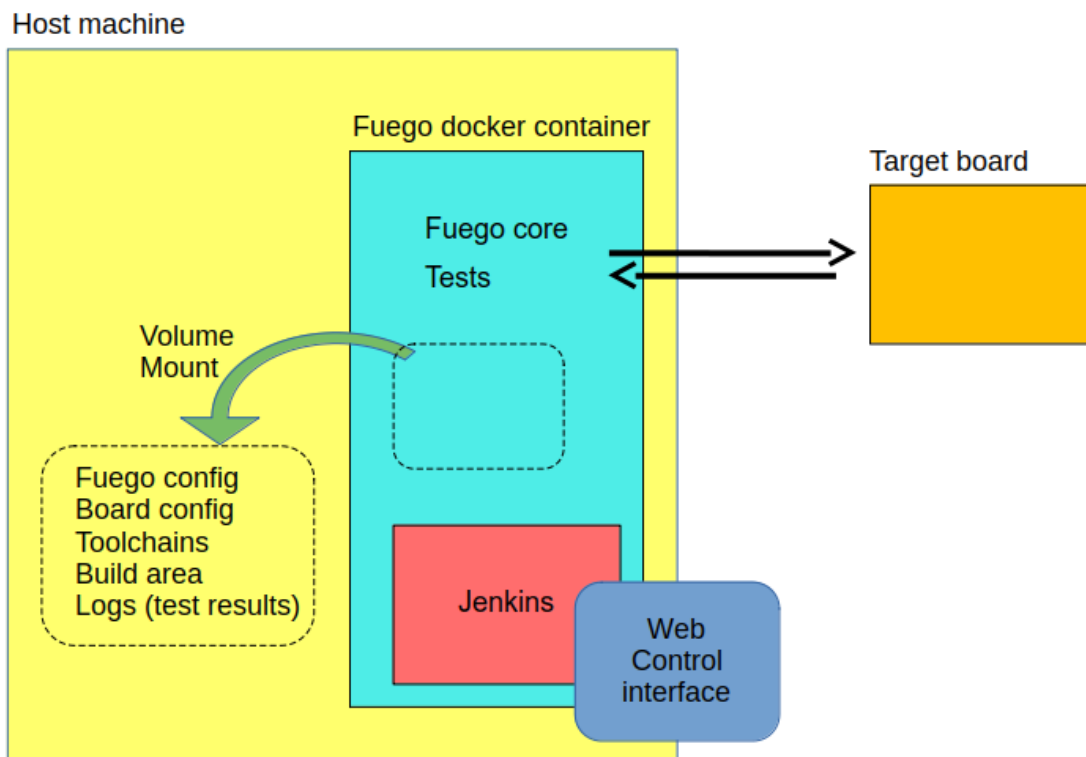
ARCHITECTURE

Fuego consists of a continuous integration system, along with some pre-packaged test programs and a shell-based test harness, running in a Docker container.:

Fuego = (Jenkins + core scripts + pre-packaged tests)
inside a container

Here's a diagram with an overview of Fuego elements:

Architecture Diagram



2.1 Major elements

The major elements in the Fuego architecture are:

- Host system
 - Fuego container instance
 - * Jenkins continuous integration system
 - Web-based user interface (web server on port 8090)
 - Jenkins Plugins
 - * Fuego core system
 - * Fuego tests
 - * Build environment
 - * Test results directory
- Target board(s)
- Fuego command line tool (not shown in the diagram above)
- Web client, for interacting with the system (not provided and not shown)

2.2 Container

By default, Fuego runs inside a Docker container, which is created when Fuego is installed. Using a container makes it is easier to install and use Fuego on a variety of different Linux distributions. The libraries and packages that Fuego needs are included in the docker container, and don't need to be installed separately. These packages have fixed versions, and don't interfere with versions of the same packages that may already be present on the host system where Fuego is installed.

2.3 Jenkins

The main user interface for Fuego is provided by the Jenkins continuous integration system.

Jenkins is used in Fuego to launch test jobs and view test results. Usually, Fuego tests are started by Jenkins automatically, in response to changes in the software. However, Jenkins can be configured to launch test jobs based on a variety of triggers, including when a user manually schedules a test to run, or at periodic intervals.

Jenkins is too big a system to describe in detail here, but it has many features and is very popular. It has an ecosystem of plugins for many kinds of extended functionality, such as integration with different source code management systems, results plotting, e-mail notifications of regressions, and more.

Fuego installs several plugins that are used by various aspects of the system. You can install additional plugins to Jenkins to suit your own requirements.

In Fuego, Jenkins is used to:

- Start tests
- Schedule tests for automatic execution
- Shows test results (particularly over time)
- Flag regressions in test results

When using Jenkins with Fuego, the Fuego administrator will add Jenkins objects (nodes and jobs) that are representative of Fuego objects. Once these objects are defined, Jenkins can then start a Fuego job. It does this by calling Fuego's `ftc` command.

The interface between Jenkins and the Fuego system is documented at [Core interfaces](#).

2.4 Pre-packaged tests

Fuego contains over 100 pre-packaged tests, ready for you to start testing “out-of-the-box”. There are tests of individual programs or features, such as ‘iputils’ or ‘pmqtest’, as well as several benchmarks in the areas of CPU performance, networking, graphics, and realtime. Fuego also includes some full test suites, like LTP (Linux Test Project). Finally, Fuego includes a set of selftests, to validate board operation or core Fuego functionality.

2.5 Fuego test definition

Fuego defines a test using a collection of files, which handle different aspects of test definition and execution.

The major elements of a Fuego test are:

- a test definition file, which has information about the test
- a “base script”, which manages test execution (`fuego_test.sh`)
- a “spec” file, that contains information about test variants
- a test program, to perform the actual test
- a parser, to convert test program output to individual testcase results

Some other files that a test might include are:

- a criteria file, for evaluating test results
- a chart config file, that controls which test results are output and in what format, in the Jenkins user interface

2.5.1 Base script vs test program

Fuego tests often provide a build system and host/target wrapper for existing test programs (such as Dhrystone, dbench or cyclictst).

In cases like this, part of the Fuego code runs on the host system, and part runs on the device under test. More specifically, the base script (`fuego_test.sh`) is run on the host machine, inside the Fuego docker container. The test program (e.g. the actual dhrystone executable) is run on the target board.

Because of this setup, it can be confusing what the words “test”, or “test script”, “test program”, or the name of the test (e.g. Dhrystone) refers to.

In order to avoid confusion, this documentation refers to the software that runs on the target board (or “device under test”) as the “test program” (always including the word ‘program’). The documentation uses the full name of the test (such as ‘Benchmark.Dhrystone’), to refer to the full set of materials used by Fuego to define a test. And it uses the term “base script”, to refer specifically to the ‘`fuego_test.sh`’ script inside a test.

As an example, for Fuego’s `Benchmark.Dhrystone` test, the following nomenclature would be used:

- ‘test’ or ‘dhrystone test’, or ‘Benchmark.Dhrystone’ = the full set of files that comprise the Fuego Dhrystone test.
- ‘base script’ = `fuego_test.sh`

- ‘test program’, or ‘dhrystone program’ = the dhrystone executable

Note: This documentation uses the terminology “test program” to refer to the software that is executed on the target board. However, it should be noted that while the test program is often a compiled program, it could be an interpreted script (such as a shell script). The phrase “test program” in this documentation does not imply that the program is always a binary object.

Not every Fuego test includes a ‘test program’. Some Fuego tests execute commands on the target board directly from the base script (running on the host) without placing any separate or additional program on the board.

2.6 Board

Fuego performs testing using a host/target configuration for building, deploying and executing tests.

Fuego executes tests on physical hardware that is accessible from the host on which Fuego is installed. The physical hardware being tested is referred to as the “target board”, or the “device under test”.

During Fuego installation, the Fuego administrator configures how Fuego accesses and controls the board by creating a board configuration file. Fuego can be configured with an arbitrary number of boards on which to run tests. As a special case, the administrator can also configure Fuego to treat the host system as a board (for self-testing).

2.6.1 Board requirements

The board might be physically connected to the host (e.g. by a serial or USB cable) or not. Fuego requires very little on the target board. It only requires that the target board have a POSIX shell and a few system utilities, as well as the capability to copy files to and from the board, and the ability to remotely execute commands on the board.

Many embedded Linux devices can satisfy these requirements with just the ‘busybox’ program and a serial or ssh connection.

Many Linux test systems assume that the system-under-test is a full desktop or server system, with sufficient horsepower to build tests and run them locally. Fuego assumes the opposite - that embedded targets may be underpowered and may not have the utilities and tools needed for building and executing tests.

2.7 Fuego core

The Fuego core consists of shell scripts (including `main.sh`) and python code that loads the data and functions for the test. Fuego also provides a command line tool, called `ftc` that is used to perform administration and management functions.

2.7.1 Test Functions

Fuego provides a library of functions (in the form of shell script code), that are used by a Fuego test to perform test operations in a way that is independent of the architecture, physical connection, or Linux distribution of the device under test.

Some of the operations that can be performed by these functions are:

- Building test programs from source
- Copying files to and from the target board

- Deploying test programs to the target board (installing them)
- Executing the test programs
- Reading the test log
- Parsing the log to determine pass or fail conditions for tests
- Parsing the log for results to display in charts

By using this library of functions, Fuego tests are insulated from the different hardware and access methods used to manage a board in a particular test lab. For example, the Fuego core and the base script for a Fuego test do not have to “know” whether a board is controlled via serial console, ssh, or some other target agent. These are configured via the board configuration file and Fuego overlay system, such that the tests themselves are independent of these details.

2.8 Different objects in Fuego

It is useful to give an overview of the major objects used in Fuego, as they will be referenced in this documentation:

Fuego core objects:

- board - a description of the device under test
- test - materials for conducting a test
- spec - one or more sets of variables for describing a test variant
- run - the results from a individual execution of a test on a board

Jenkins objects:

- node - the Jenkins object corresponding to a Fuego board
- job - a Jenkins object corresponding to a combination of board, spec, and test
- build - the test results, from Jenkins perspective. This corresponds to a Fuego ‘run’

Fuego consists of both a front-end and a back-end. To avoid confusion, different names are used to describe the front-end and back-end objects used by the system. Jenkins is the front-end, and the Jenkins objects have rough counterparts in the Fuego core, as follows:

Jenkins object	Corresponding Fuego object
node	board
job	test
build	run

2.9 Jenkins operations

This section explains how Jenkins works as part of Fuego.

- When a test job is initiated, Jenkins starts a slave process to run the test that corresponds to that job
- The slave (slave.jar) runs a small shell script fragment, that is specified in the configuration (config.xml) for the job
 - This script runs the `ftc run-test` command.
 - * `ftc` executes the Fuego core that does the actual building, deploying and execution of a test.
 - * Details of the core execution of the test are described below in the *Test execution* section.

- While a test is running, Jenkins accumulates the log output from the test execution, and displays it to the user (if they are watching the console log)
- Jenkins provides a web UI for browsing the nodes, jobs, and test results (builds), and displaying graphs for benchmark data.

By default, Jenkins is installed as part of the Fuego system. However, it is possible to use Fuego without using Jenkins, by calling the Fuego command line tool (`ftc`) directly from your own testing infrastructure or CI system (e.g. `gitlab`).

2.10 Test execution

This section describes the major elements and operations of the Fuego core when a test is executed.

When a test is started, Fuego generates a test environment, consisting of variables and functions from the core system, using something called the overlay generator. The test environment is placed into a file called `prolog.sh` and loaded into the currently running shell environment.

Details of the test environment generation are described below.

Each Fuego test has a base script, called `fuego_test.sh`, that defines a few functions with operations that are specific to that test. The base script is also loaded into the currently running shell environment.

The Fuego core performs the test using variables and calling functions from: 1) the Fuego core, 2) `prolog.sh` (the test environment), and 3) `fuego_test.sh` (the base script).

A Fuego test is executed in a series of phases which perform different operations during the test. The most critical operation is running the actual test program on the board. Specifically, the Fuego core calls the base script's `test_run` function, which executes the test program on the board and collects its output.

The Fuego core also collects additional information from the board, and cleans up after the test. Finally, the Fuego core analyzes the test output, by parsing the test logs, and generates data files containing test results.

2.11 Test environment file generation

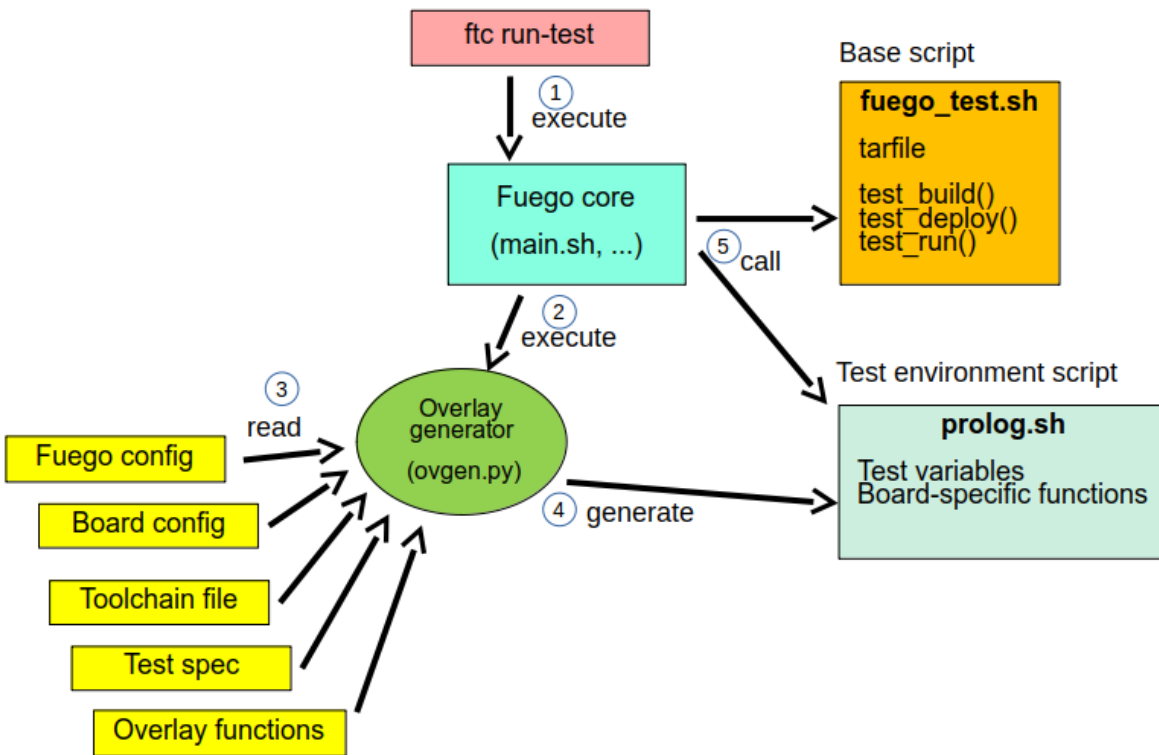
At the start of test execution, Fuego gathers information from the Fuego system and creates a test environment file.

Information from the board configuration file (and other sources) is used to create variables and functions that are specific to the current test invocation and the board under test. These are placed into the test environment file, which is called `prolog.sh` and located in the test's log directory. These items are then used during test execution.

This operation is referred to as “overlay generation”, because some of the variables and functions come from class files that can have their values overridden (or “overlay”ed) by elements of the board configuration file.

- The overlay generator takes the following as input:
 - Environment variables passed by Jenkins and `ftc`
 - The board configuration file
 - The toolchain configuration file
 - The test spec for the test
 - The overlay class files

Fuego uses the variables `TOOLCHAIN`, `DISTRIB`, `TRANSPORT`, and `BOARD_CONTROL` in the board configuration file to determine the variables and functions to include in the test environment file.



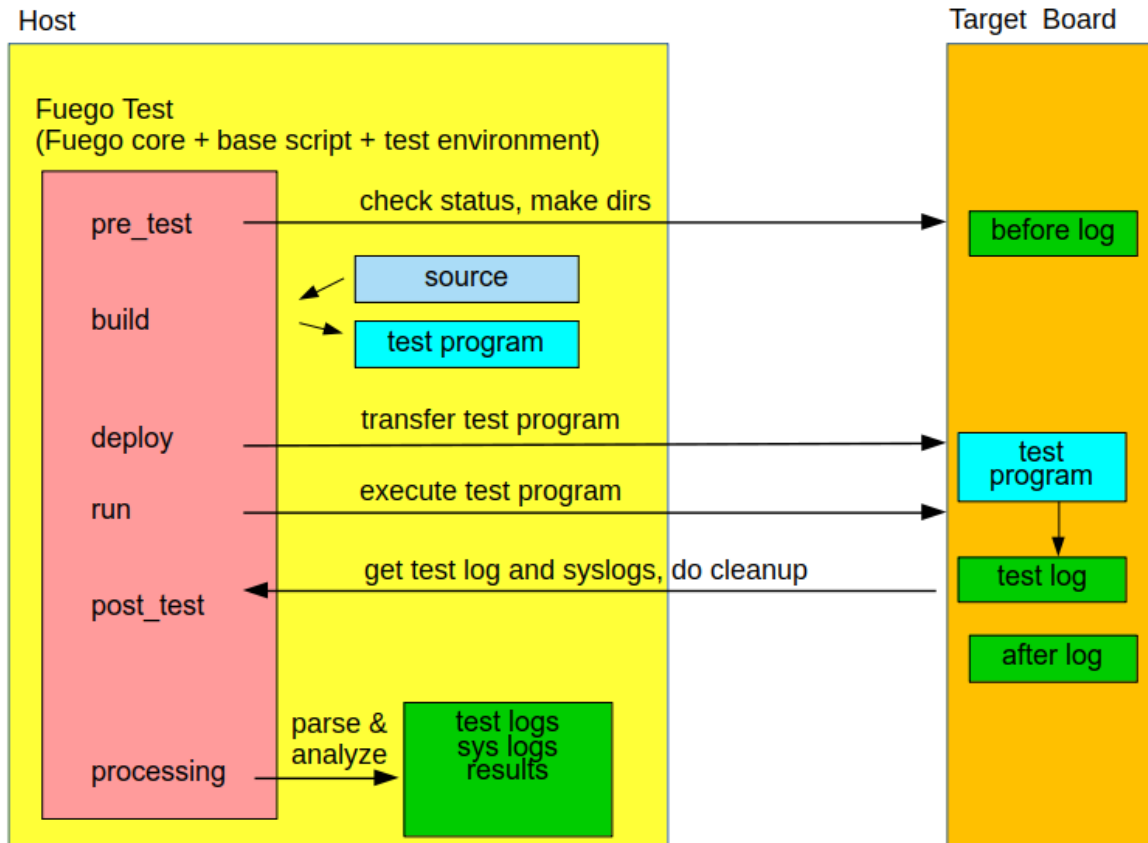
2.12 Fuego test phases

Fuego runs through a series of well-defined phases during test execution. The reason for having distinct phases is that it allows for easier debugging of test execution, during the development of a test. Some of the test phases are optional, and a user can manually control which individual phases of a test are executed. Also, the user can control which phases show extra debug information during a test.

The major test phases are:

- `pre_test`
- `build`
- `deploy`
- `run`
- `post_test`
- `processing`

Each of these are described in the sections after this diagram.



2.12.1 pre_test

The `pre_test` phase consists of making sure the target is booted and running, and preparing the workspace for the test. In this phase test directories are created, and the kernel version on the target board is recorded.

The 'before' syslog is created, and filesystems are synced and buffer caches dropped, in preparation for any filesystem tests.

If the function "test_pre_check" is defined in the base script, it is called to do any `pre_test` operations, including checking to see if required variables are set, programs or features are available on the host or target board, or whether any other test dependencies are met.

If the dependencies are not met, the test is aborted at this phase, with a result of `ERROR` and messages indicating the reason for the failure.

2.12.2 build

During this phase, the test program source is installed on the host (inside a build area in the container), and the software for the test is actually built. The toolchain specified by `TOOLCHAIN` is used to build the software.

The source code may be provided in the form of a tarball file, or it may be specified by the test as a reference to a git repository.

This phase is split into multiple parts:

- `pre_build` - build workspace is created, a build lock is acquired and the tarball is unpacked

- The `unpack` function is called during `pre_build`. This function unpacks the source tar file, if one is associated with this test
- `test_build` - the `test_build` function from `fuego_test.sh` is called
 - usually this consists of ‘`make`’, or ‘`configure ; make`’
- `post_build` - (empty for now)

2.12.3 deploy

The purpose of this phase is to copy the test programs, and any required supporting files, to the target.

This consists of 3 sub-phases:

- `pre_deploy` - `cd`’s to the build directory
- `test_deploy` - the base script’s ‘`test_deploy`’ function is called.
 - Usually this consists of tarring up needed files, copying them to the target with ‘`put`’, and then extracting them there
 - Items should be placed in the directory `$BOARD_TESTDIR/fuego.$TESTDIR` directory on the target
- `post_deploy` - removes the build lock

2.12.4 run

In this phase the test program on the target board is actually executed.

This executes the ‘`test_run`’ function defined in the base script for the test, which usually consists of one or more calls to the `report` function, which executes the test program on the target board and collects the standard out from the program. This output is saved as the testlog for the test. Note that the `report` function saves the testlog on the target. It is collected from the target later, for post-processing.

The run phase may include additional commands to prepare the system for test operation and clean up after the execution of the test program.

2.12.5 post_test

In this phase, the test log is retrieved (fetched) from the target and stored on the host. Also in this phase, the board is “cleaned up”, which means that test directories and logs are removed on the target board, and any leftover processes related to this test that are running on the board are stopped.

2.12.6 processing

In the processing phase of the test, the results from the test log are evaluated. The `test_processing` function of the base test script is called.

For functional tests:

Usually, this phase consists of one or more calls to ‘`log_compare`’, to determine if a particular string occurs in the testlog. This phase determines whether the test passed or failed, and the base test script indicates this (via it’s exit code) to the Jenkins interface.

For benchmarking tests:

This phase consists of parsing the testlog, using `parser.py`, and also running `dataload.py` to save data for plot generation.

Also, a final analysis is done on the system logs is done in this step (to detect things like Kernel Ooopses that occurred during the test).

2.12.7 phase relation to base script functions

Some of the phases are automatically performed by Fuego, and some end up calling a routine in the base script (or use data from the base script) to perform their actions. This table shows the relation between the phases and the data and routines that should be defined in the base script.

It also shows the most common commands utilized by base script functions for this phase.

phase	relationship to base script	common operations
pre_test	calls 'test_pre_check'	assert_define, is_on_target, check_process_is_running
build	uses the 'tarfile' definition, calls 'test_build'	patch, configure, make
deploy	Calls 'test_deploy'	put
run	calls 'test_run'	cmd, report, report_append
post_test	calls 'test_cleanup'	kill_procs
process- ing	calls 'test_processing'	log_compare

2.12.8 Other scripts and programs

- `parser.py` - for parsing test results
- `criteria.json` - for analyzing test results

A test might also include a file called `parser.py`. In fact, every benchmark test should have one. This file is a python module which is run to extract results and data values from the log.

This script is run inside the docker container, after a test is finished. The Fuego log parsing system loads this module as part of test processing.

A benchmark program measures some attribute of the system during a test, and produces one or more values called 'metrics'. These values are emitted by the benchmark test into the test log, and the Fuego parser retrieves these values and uses them to evaluate the pass/fail status of the benchmark. These values are saved as part of the test results, and are used by plotting software to show charts of test results in the Jenkins interface.

Tests may also include a file called `criteria.json` which is used to determine whether test results constitute a pass or fail result. For example, for benchmark tests, the system can collect a number from the test program, but it is up to the system to determine whether that number represents an acceptable value (pass), or a failure or regression (fail). The `criteria.json` file has data about metric thresholds, for benchmark tests, and about test results that can be ignored, for functional tests, to allow for automating this results processing.

2.13 Command line tool

Fuego includes a command line tool for performing administrative and management operations, and for executing tests. This command line tool is called `ftc`. Details of `ftc` commands can be found in the section [*Command Line Tool*](#).

INSTALL AND FIRST TEST

This tutorial has some short setup instructions if you just want to get a taste of what Fuego is like. This allows you to experiment with Fuego and try out some tests to see what it looks like and how it works, without investing a lot of time (well, except for the first container build).

In this configuration, we will show you how to install Fuego and run a test on a ‘docker’ board, which is the docker container where Fuego itself is running, on your host machine.

Obviously, this is not useful for testing any real hardware. It is intended only as a demonstration of Fuego functionality. For instructions to set up a real board, try the Fuego Quickstart Guide or the [Installing Fuego](#) page.

3.1 Overview

An overview of the steps is:

1. Install pre-requisite software
2. Download the Fuego repository
3. Build your Fuego container
4. Start the container
5. Add the ‘docker’ board to Jenkins
6. Add some sample tests
7. Access the Jenkins interface
8. Run a test

These steps are described below.

3.2 Step details

To install and run Fuego, you need to have git and docker installed on your system.

On Ubuntu, try the following commands:

```
$ sudo apt install git docker.io
```

To download Fuego, and build and start the container, type the following commands at a Linux shell prompt:

```
$ git clone https://bitbucket.org/fuegotest/fuego.git
$ cd fuego
$ ./install.sh
$ ./start.sh
```

The third step (with `command:./install.sh`) will take some time - about 45 minutes on an average Linux machine. You might want to go make yourself a sandwich, (or go watch the [Fuego introduction video](#)). This step is building the “Fuego” distribution of Linux (based on Debian) and putting it into the Fuego docker container. You will also need a connection to the Internet with fairly decent bandwidth.

When you run the ‘start.sh’ script, the terminal will be placed at a shell prompt, as the root user inside the docker container. The container will run until you exit this shell. You should leave it running for the duration of your testing.

The next steps populate the Jenkins system objects used for testing:

At the shell prompt inside the container type the following:

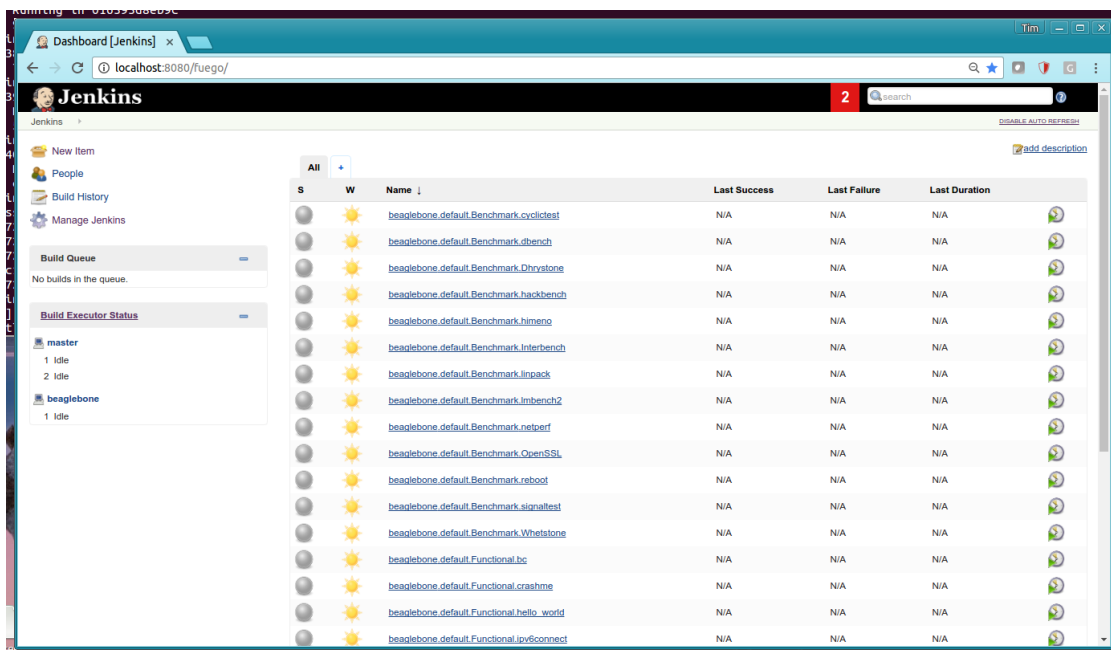
```
# ftc add-node -b docker
# ftc add-jobs -b docker -t Functional.batch_smoketest
```

This will add the ‘docker’ node in the Jenkins interface, and a small set of tests. The ‘docker’ node simulates an embedded “board” using the Fuego docker container itself. The “smoketest” batch test has about 20 tests that exercise a variety of features in a Linux system. After running these commands, a set of jobs will appear in the Jenkins interface.

```
$ firefox http://localhost:8090/fuego
```

To access the Fuego interface (Jenkins) you can use any browser - not just Firefox. By default the Fuego interface runs on your host machine, on port 8090, with URL path “/fuego”.

In your browser, you should see a screen similar to the following:



3.3 Run a test

3.3.1 Manually start a test

To run a job manually, do the following:

- Go to the Jenkins dashboard (in the main Jenkins web page),
- Select the job (which includes the board name and the test name)
- Click “Build job” (Jenkins refers to running a test as “building” it.)

A few very simple jobs you might start with are:

- Functional.hello_world
- Benchmark.Drhystone

You can also start a test manually by clicking on the circle with a green triangle, on the far right of the line with the job name, in the Jenkins dashboard.

When you run a test, the test software is built from source, sent to the machine (in this case the Fuego docker container), and executed. Then the results are collected, analyzed, and displayed in the Jenkins interface.

3.3.2 Look at the results

When the test has completed, the status will be shown by a colored ball by the side of the test in the dashboard. Green means success, red means failure, and grey means the test did not complete (it was not run or it was aborted).

You can get details about the test by clicking on the links in the history list. You can see the test log (the output from the test program), by clicking on the “testlog” link. You can see the steps Fuego took to execute the test by clicking on the “console log” link on the job page. And you can see the formatted results for a job, and job details (like start time, test information, board information, and results) in the test’s ‘run.json’ file.

3.4 Conclusions

Hooray! You have completed your first Fuego test. Although this was not a test on real hardware, you (and Fuego) have completed a lot of stuff behind the scenes. You have:

- Downloaded your own distribution of Fuego and installed it in a docker container
- Added a fake “board” to Jenkins
- Added tests to run on the board
- Executed a test

During this last step, Fuego did the following:

- Built the test program from source code
- Downloaded the test to the “board”
- Executed the test
- Retrieved the test log from the board
- Analyzed the log, and formatted results for display in Jenkins

Whew! That’s a lot of work. And all you had to do (after initial installation) was click a button.

3.5 What do do next?

In order to use Fuego in a real Continuous Integration loop, you need to do a few things:

- Configure Fuego to work with your own board or product
- Select the set of tests you would like to run on your board
- Customize benchmark thresholds and functional baselines for those tests, for your board
- Configure Fuego jobs to be triggered after new software is installed on the board

Fuego does not currently have support for installing new system software (the kernel and root filesystem) on boards itself. This is something you need to automate outside of Fuego, if you plan to use Fuego in your CI loop for system software.

Usually, Fuego users create their own Jenkins job which provisions the board (installs the kernel and/or root filesystem for their chosen Linux distribution), and then triggers Fuego jobs, after the new software is installed on the board.

See further instructions see the Fuego *Quickstart Guide*, *Adding a Board*, *Adding a toolchain* or the *Installing Fuego* page.

FUEGO QUICKSTART GUIDE

Running tests from Fuego on your hardware can be accomplished in a few simple steps.

Note: This is the Quickstart Guide. More detailed instructions can be found at [Installing Fuego](#).

4.1 Overview

The overview of the steps is:

1. install pre-requisite software
2. download the fuego repository
3. build your fuego container
4. start the container
5. access the interface
6. add your board to fuego
7. run a test

These steps are described below.

4.2 Install pre-requisite software

To retrieve the fuego software and create the docker image for it, you need to have git and docker installed on your system.

On Ubuntu, try the following commands::

```
$ sudo apt install git docker.io
```

4.3 Download, build, start and access

To accomplish the last 6 steps, do the following from a Linux command prompt::

```
$ git clone https://bitbucket.org/fuegotest/fuego.git
$ cd fuego
$ ./install.sh
$ ./start.sh
$ firefox http://localhost:8090/fuego
```

The fourth step (with `./install.sh`) will take some time - about 45 minutes on my machine. This is the main step that builds the Fuego docker container.

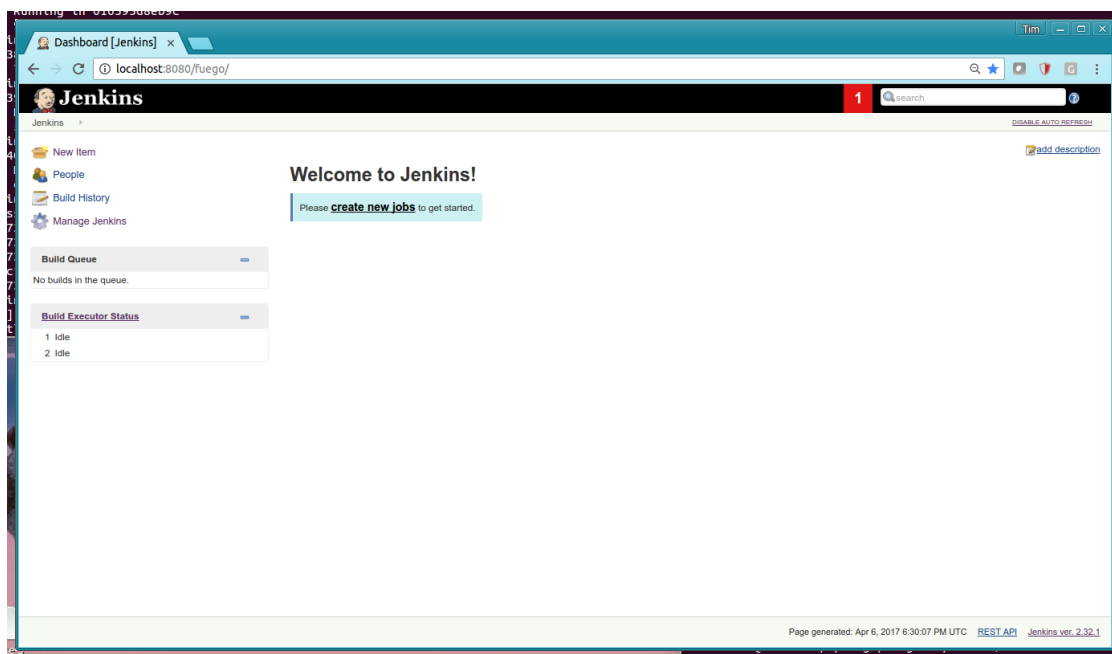
When you run the `'start.sh'` script, the terminal where this is run will be placed at a shell prompt, as the root user, inside the docker container. The container will run until you exit this shell. You should leave it running for the duration of your testing.

Note: If you are experimenting with the unreleased version of Fuego in the `'next'` branch, then please replace the `'git clone'` command in the instructions above with this (i.e. add `'-b next'`):

```
$ git clone -b next https://bitbucket.org/fuegotest/fuego.git
```

On the last step, to access the Fuego interface you can use any browser - not just Firefox. By default the Fuego interface runs on your host machine, on port 8090, with URL path `"/fuego"`.

In your browser, you should see a screen similar to the following:



We will now add items to Fuego (and this screen) so you can begin testing.

4.4 Add your board to fuego

To add your own board to Fuego, there are three main steps:

- 1. create a test directory on the target
- 2. create a board file (on the host)
- 3. add your board to the Jenkins interface

You can find detailed instructions for adding a board at: [Adding a Board](#)

However, here is a quick list of steps you can do to add a your own board, and a sample ‘docker’ board to Fuego:

4.4.1 Create a test directory on your board

Login to your board, and create a directory to use for testing::

```
$ ssh root@your_board
<board>$ mkdir /home/a
<board>$ exit
```

If not using ssh, use whatever method you normally use to access the board.

4.4.2 Create board file

Now, create your board file. The board file resides in <fuego-dir>/fuego-ro/boards, and has a filename with the name of the board, with the extension “.board”.

Do the following:

```
$ cd fuego-ro/boards
$ cp template-dev.board myboard.board
$ vi myboard.board
```

Edit the variables in the board file to match your board. Most variables can be left alone, but you will need to change the IPADDR, TOOLCHAIN and ARCHITECTURE variables, and set the BOARD_TESTDIR to the directory you just created above.

For other variables in the board file, or specifically to use a different transport than SSH, see more complete instructions at: [Adding a Board](#)

4.4.3 Add boards to the Jenkins interface

Finally, add the board in the Jenkins interface.

In the Jenkins interface, boards are referred to as “Nodes”.

At the container shell prompt, run the following command:

- (container prompt)\$ ftc add-nodes -b myboard docker

This will add your board as a node, as well as a ‘docker’ node in the Jenkins interface.

4.5 Install a toolchain

If you just wish to run experiment with Fuego, without installing your own board, you can use the existing ‘docker’ board. This will run the tests inside the docker container on your host machine. This requires little setup, and is intended to let people try Fuego to see how the interface and tests work, without having to set up their own board.

If you are running an ARM board with a Debian-based distribution on it, you can install the Debian ARM cross-compilers into the docker container with the following command (inside the container):

- (container prompt)\$ /fuego-ro/toolchains/install_armhf_toolchain.sh

If you are installing a some other kind of board (different architecture, different root filesystem layout, or different shared library set), you will need to install a toolchain for your board inside the docker container.

Please follow the instructions at: [Adding a toolchain](#) to do this.

4.6 Now select some tests

In order to execute tests using the Jenkins interface, you need to create Jenkins “jobs” for them. You can do this using the ‘ftc add-jobs’ command.

These commands are also executed at the shell prompt in the docker container.

You can add jobs individually, or you can add a set of jobs all at once based on something called a ‘testplan’. A testplan is a list of Fuego tests with some options for each one. You can see the list of testplans in your system with the following command:

- (container prompt)\$ ftc list-plans

To create a set of jobs for the ‘docker’ board on the system, do the following:

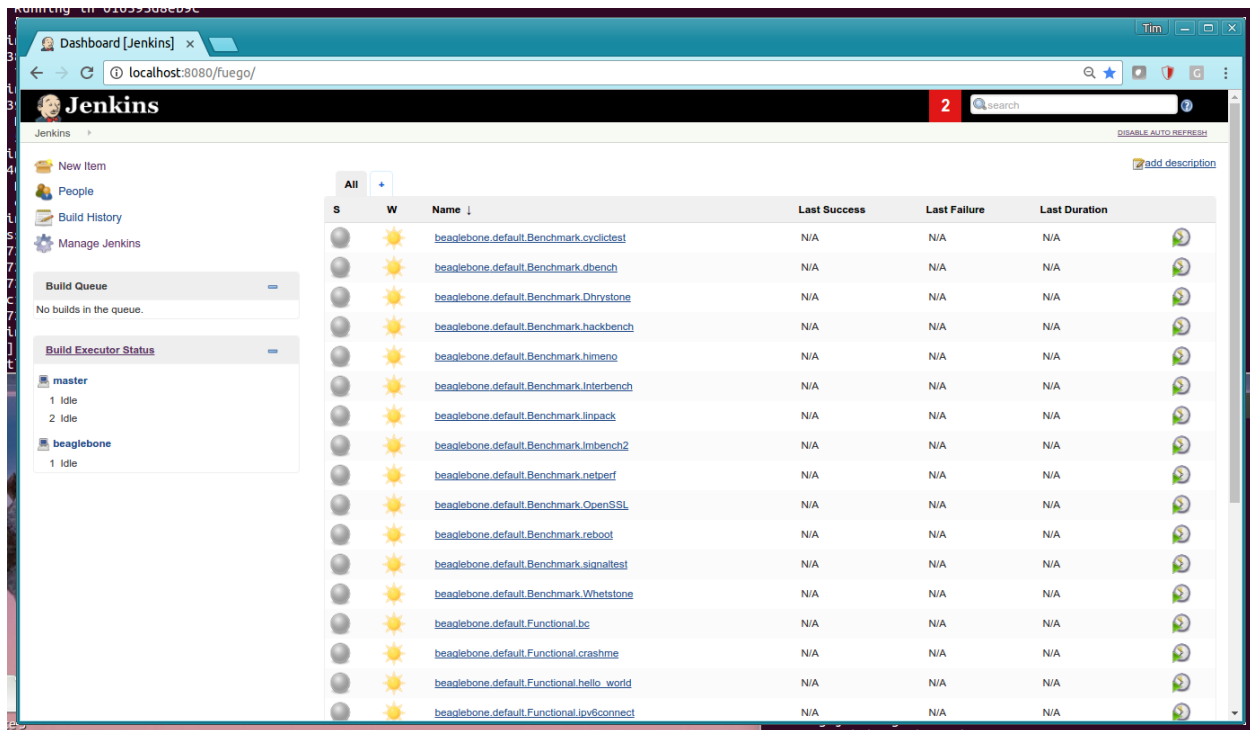
- (container prompt)\$ ftc add-jobs -b docker -p testplan_docker

To create a set of jobs for your own board (assuming you called it ‘myboard’), do the following:

- (container prompt)\$ ftc add-jobs -b myboard -p testplan_smoketest

The “smoketest” testplan has about 20 tests that exercise a variety of features in a Linux system. After running these commands, a set of jobs will appear in the Jenkins interface.

Once this is done, your Jenkins interface should look something like this:



4.7 Run a test

To run a job manually, you can do the following:

- Go to the Jenkins dashboard (in the main Jenkins web page),
- Select the job (which includes the board name and the test name)
- Click “Build job” (Jenkins refers to running a test as “building” it.)

You can also click on the circle with a green triangle, on the far right of the line with the job name, in the Jenkins dashboard.

When the test has completed, the status will be shown by a colored ball by the side of the test in the dashboard. Blue means success, red means failure, and grey means the test did not complete (was not run or was aborted). You can get details about the test run by clicking on the link in the history list.

4.8 Additional Notes

4.8.1 Other variables in the board file

Depending on the test you want to run, you may need to define some other variables that are specific to your board or the configuration of the filesystem on it. Please see Adding a Board for detailed instructions and a full list of variables that may be used on the target.

4.8.2 the Jenkins interface

See Jenkins User Interface for more screenshots of the Jenkins web interface. This will help familiarize you with some of the features of Jenkins, if you are new to using this tool.

4.9 Troubleshooting

If you have problems installing or using Fuego, please see our Troubleshooting Guide

INSTALLING FUEGO

This page describes the steps to install Fuego on your Linux machine. It includes detailed descriptions of the operations, for both users and developers.

Tip: If you are interested in a quick outline of steps, please see the *Fuego Quickstart Guide* instead.

5.1 Overview

The overview of the steps is:

1. Install pre-requisite software
2. Download the Fuego repository
3. Build your Fuego container
4. Start the container
5. Access the Jenkins interface

5.2 Install pre-requisite software

To retrieve the Fuego software and create the Docker image for it, you need to have **git** and **Docker** installed on your system.

On Ubuntu, try the following commands:

```
$ sudo apt-get install git
$ sudo apt-get install docker.io
```

These commands may be different for other distributions of Linux (such as Fedora, RedHat, CentOS, Mint, etc.)

5.3 Overview of remaining steps

Steps 2 through 5 of the installation can be performed with the following Linux commands:

```
$ git clone https://bitbucket.org/fuegotest/fuego.git
$ cd fuego
$ ./install.sh
$ ./start.sh
$ firefox http://localhost:8090/fuego
```

These steps and commands will be described in the sections that follow.

5.4 Install the Fuego repositories

The Fuego system is contained in 2 git repositories. One repository is called `fuego` and the other is called `fuego-core`. The `fuego-core` repository is installed inside the `fuego` directory, at the top level of that repository's directory structure. This is done automatically during the install of Fuego. You do not need to clone the `fuego-core` repository manually yourself.

The reason to have two repositories is that they hold different pieces of the Fuego system, and this allows for them to be upgraded independently of each other.

The repositories are hosted on `bitbucket.org`, under the the `fuegotest` account.

5.4.1 Fuego repository

The `fuego` repository has the code and files used to build the Fuego docker container. It also has the `fuego-ro` directory, which has board definition files, various configuration files, miscellaneous scripts, and other items which are used by Fuego for container management or other purposes.

5.4.2 Fuego-core repository

The `fuego-core` repository has the code which implements the core of the Fuego test execution engine, as well as the pre-packaged tests included with the system. This includes the overlay generator, the results parser, the Fuego shell function library, the directory of tests, and the main Fuego command line tool `ftc`.

5.4.3 Downloading the repository

You can use `git clone` to download the main Fuego repository, like so:

```
$ git clone https://bitbucket.org/fuegotest/fuego.git
$ cd fuego
```

After downloading the repositories, switch to the `fuego` directory, as shown in the example.

Note that these git commands will download the 'master' branch of the repository, which is the current main released version of Fuego.

Downloading a different branch

If you are experimenting with an unreleased version of Fuego in the ‘next’ branch, then please replace the ‘git clone’ command in the instructions above with these:

```
$ git clone -b next https://bitbucket.org/fuegotest/fuego.git
$ cd fuego
```

This uses `-b next` to indicate a different branch to check out during the clone operation.

5.5 Create the Fuego container

The third step of the installation is to run `install.sh` to create the Fuego docker container. While in the `fuego` directory, run the script from the current directory, like so:

```
$ ./install.sh
```

`install.sh` uses Docker and the Dockerfile in the `fuego` directory to create a Docker container with the Fuego Linux distribution.

This operation may take a long time. It takes about 45 minutes on my machine. This step assembles a nearly complete distribution of Linux, from binary packages obtained from the Internet.

This step requires Internet access. You need to make sure that you have proxy access to the Internet if you are behind a corporate firewall.

Please see the section “Alternative Installation Configuratons” below for other arguments to `install.sh`, or for alternative installation scripts.

5.5.1 Fuego Linux distribution

The Fuego Linux distribution is a distribution of Linux based on Debian Linux, with many additional packages and tools installed. These additional packages and tools are required for aspects of Fuego operation, and to support host-side processes and services needed by the tests included with Fuego.

For example, the Fuego distribution includes:

- the **Jenkins** continuous integration server
- the `netperf` server, for testing network performance.
- the `ttc` command, which is a tool for board farm management
- the `python jenkins` module, for interacting with Fuego’s Jenkins instance
- and many other tools, programs and modules used by Fuego and its tests

Fuego commands execute inside the Fuego docker container, and Fuego operations initiate in the container, but may access hardware (such as USB ports, networking, or serial ports) that are outside the container.

5.5.2 Configuring for ‘privileged’ hardware access

In many configurations, Fuego can perform its operations using only network operations. However, depending on the configuration of your boards, or your lab hardware, and the relationship between your host and target computers used for testing, you may need to access other hardware on your host machine.

To do that, you can create a ‘privileged’ Fuego container, using the `--priv` options with `install.sh`:

```
$ ./install.sh --priv
```

Customizing the privileged container

Note that using `--priv` causes `install.sh` to use a different container creation script. Normally (in the non `--priv` case), `install.sh` uses `fuego-host-scripts/docker-create-container.sh`. When `--priv` is used, Fuego uses `fuego-host-scripts/docker-create-usb-privileged-container.sh`.

This latter script (`docker-create-usb-privileged-container.sh`) can be edited, before running `install.sh`, to change the set of hardware devices that the Docker container will have privileged access to.

This is done by adding more bind mount options to the `docker create` command inside this script. Explaining exactly how to do this is outside the scope of this documentation. Please see documentation and online resources for the Docker system for information about this.

The script currently creates bind mounts for:

- `/dev/bus/usb` - USB ports, and newly created ports
- `/dev/ttyACM0` - serial port 0
- `/dev/ttyACM1` - serial port 1
- `/dev/serial` - general serial ports, and newly created ports

If you experience problems with Fuego accessing hardware on your host system, you may need to build the Fuego docker container using additional bind mounts that are specific to your configuration. Do so by editing `docker-create-usb-privileged-container.sh`, removing the old container, and re-running `./install.sh --priv` to build a new container with the desired privileges.

5.5.3 Using an different container name

By default, `install.sh` creates a Docker image called `fuego` and a Docker container called `fuego-container`. There are some situations where it is desirable to use different names. For example, having different container names is useful for Fuego self-testing. It can also be used to do A/B testing when migrating from one release of Fuego to the next.

You can provide a different name for the Fuego image and container, by supplying one on the command line for `install.sh`, like so:

```
$ ./install.sh my-fuego
```

This would create a Docker image named `my-fuego` and a Docker container named `my-fuego-container`

5.6 Start the Fuego container

To start the Fuego docker container, use the `start.sh` script.

```
$ ./start.sh
```

5.6.1 Using a different container name

By default, `start.sh` will start the container named `fuego-container` (which is the default Fuego docker container name). However, if you created a different container name, you can specify the name on the command line, like so:

```
$ ./start.sh my-fuego-container
```

When you run the `start.sh`, the terminal where the script is run will be placed at a shell prompt inside the Docker container. The session will be logged in as the root user inside the container. The container will run until you exit this top-level shell. Therefore, you should leave it (the shell and the terminal that your ran `start.sh` from) running for the duration of your testing.

5.7 Access the Fuego Jenkins web interface

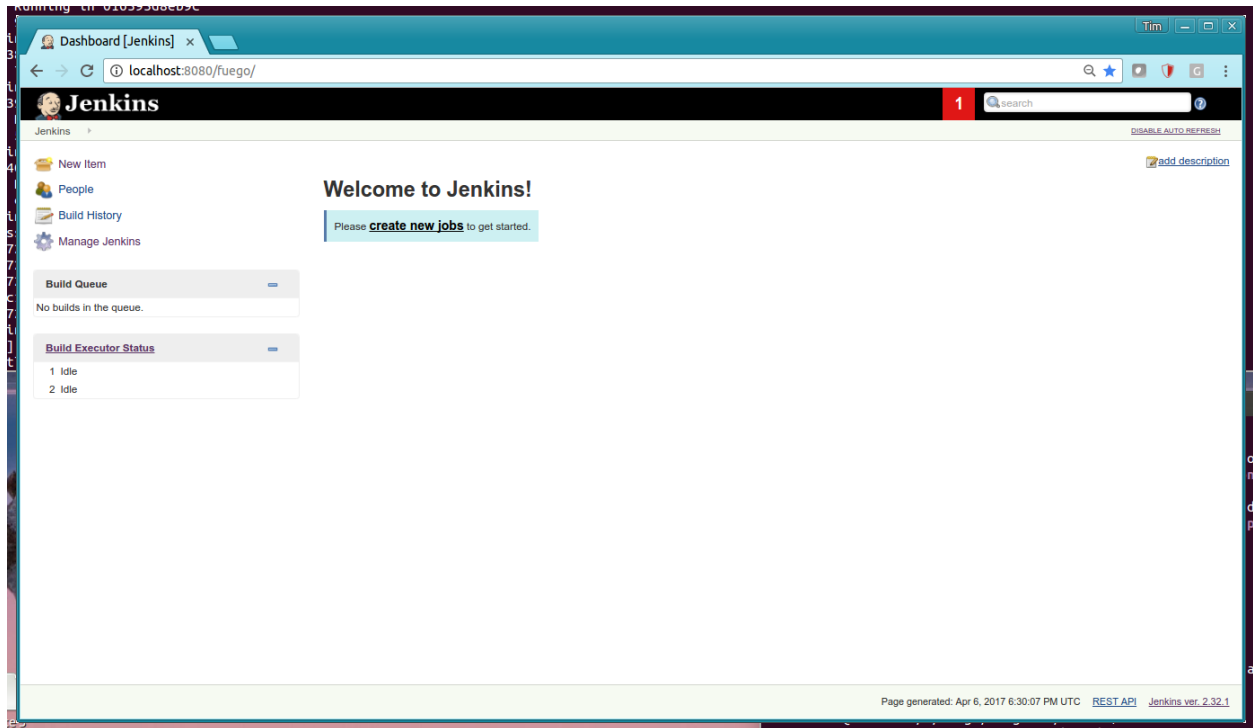
Fuego includes a version of Jenkins and a set of plugins as part of its system. Jenkins is running inside the Fuego docker container. By default the Fuego Jenkins interface runs on port `8090`, with an URL path `/fuego`.

Here is an example showing use of firefox to access the Jenkins interface with Fuego

```
$ firefox http://localhost:8090/fuego
```

To access the Fuego interface you can use any browser - not just Firefox.

In your browser, you should see a screen similar to the following:



Note that this web interface is available from any machine that has access to your host machine via the network. This means that test operations and test results are available to anyone with access to your machine. You can configure Jenkins with different security to avoid this.

5.8 Access the Fuego docker command line

For some Fuego operations, it is handy to use the command line (shell prompt) inside the Docker container. In particular, parts of the remaining setup of your Fuego system involve running the `ftc` command line tool.

Some `ftc` commands can be run outside the container, but other require that you execute the command inside the container.

To start another shell prompt inside the currently running Fuego docker container, you can use the script `fuegosh`. This helper script is located at: `fuego/fuego-ro/scripts/fuegosh`. You may find it convenient to copy this script to a `bin` directory on your system (either `/usr/local/bin` or `~/bin`) that is already in your `PATH`.

If you run `fuegosh`, it will place you at a shell prompt inside the Fuego docker container, like so:

```
$ fuegosh
root@hostname:/#
```

5.9 Remaining steps

Fuego is now installed and ready for test operations. However, some steps remain in order to use it with your hardware. You need to:

- Add one or more hardware boards (board definition files)
- Add a toolchain
- Populate the Jenkins interface with test jobs

These steps are described in subsequent sections of this documentation.

See the following sections for more information:

- *[Adding a Board](#)*
- *[Adding a Toolchain](#)*
- *[Adding Test Jobs to Jenkins](#)*

5.10 Alternative installation configurations

The default installation of Fuego installs the entire Fuego system, including Jenkins and the Fuego core, into a Docker container running on a host system, which Jenkins running on port 8090. However, it is possible to install Fuego in other configurations.

The configuration alternatives that are supported are:

- Install using a different TCP/IP port for Jenkins
- Install without the Jenkins server
- Install directly to your host (not inside a container)

5.10.1 With a different Jenkins TCP/IP port

By default the Fuego uses TCP/IP port 8090, but this can be changed to another port. This can be used to avoid a conflict with a service already using port 8090 on your host machine, or so that multiple instances of Fuego can be run simultaneously.

To use a different port than 8090 for Jenkins, specify it after the image name on the command line when you run `install.sh`. Note that this means that you must specify a Docker image name in order to specify a non-default port. For example:

```
$ ./install.sh fuego 7777
```

This would install Fuego, with an Docker image name of `fuego`, a Docker container name of `fuego-container`, and with Jenkins configured to run on port 7777

5.10.2 Without Jenkins

Some Fuego users have their own front-ends or back-ends, and don't need to use the Jenkins CI server to control Fuego tests, or visualize Fuego test results. `install.sh` supports the option `--nojenkins` which produces a Docker container without the Jenkins server. This reduces the overhead of the Docker container by quite a bit, for those users.

Inside the Docker container, the Fuego core is still available. Boards, toolchains, and tests are configured normally, but the `ftc` command line tool is used to execute tests. There is no need to use any of the `ftc` functions to manage nodes, jobs or views in the Jenkins system. `ftc` is used to directly execute tests using `ftc run-test`, and results can be queried using `ftc list-runs` and `ftc gen-report`.

When using Fuego with a different results visualization backend, the user will use `ftc put-run` to send the test result data to the configured back end.

5.10.3 Without a container

Usually, for security and test reproducibility reasons, Fuego is executed inside a Docker container on your host machine. That is, the default installation of Fuego will create a Docker container using all the software that is needed for Fuego's tests. However, in some configurations it is desirable to execute Fuego directly on a host machine (not inside a Docker container). A user may have a dedicated machine, or they may want to avoid the overhead of running a Docker container.

A separate install script, called `install-debian.sh` can be used in place of `install.sh` to install the Fuego system onto a Debian-based Linux distribution.

Please note that installing without a container is not advised unless you know exactly what you are doing. In this configuration, Fuego will not be able to manage host-side test dependencies for you correctly.

Please note also that executing without a container presents a possible security risk for your host. Fuego tests can run arbitrary bash instruction sequences as part of their execution. So there is a danger when running tests from unknown third parties that they will execute something on your test host that breaches the security, or that inadvertently damages your filesystem or data.

However, despite these drawbacks, there are test scenarios (such as installing Fuego directly to a target board), where this configuration makes sense.

ADDING A BOARD

6.1 Overview

To add your own board to Fuego, there are five main steps:

1. Make sure you can access the target via ssh, serial or some other connection
2. Decide whether to use an existing user account, or to create a user account specifically for testing
3. Create a test directory on the target
4. Create a board file (on the host)
5. Add your board as a node in the Jenkins interface

6.1.1 1 - Set up communication to the target board

In order for Fuego to test a board, it needs to communicate with it from the host machine where Fuego is running.

The most common way to do this is to use ‘ssh’ access over a network connection. The target board needs to run an ssh server, and the host machine connects to it using the ‘ssh’ client.

The method of setting an ssh server up on a board varies from system to system, but sample instructions for setting up an ssh server on a raspberry pi are located here: [*Raspberry Pi Fuego Setup*](#)

Another method that can work is to use a serial connection between the host and the board’s serial console. Setting this up is outside the scope of this current documentation, but Fuego uses the “serio” package to accomplish this. I

6.1.2 2 - Decide on user account for testing (creating one if needed)

On your target board, a user account is required in order to run tests.

The user account used by Fuego is determined by your board file, which you will configure manually in step 4. You need to decide which account to use. There are three options:

- use the root account
- use an existing account
- use a new account, dedicated to testing

There are pros and cons to each approach.

My personal preference is to use the root account. Several tests in Fuego require root privileges. If you are working with a test board, that you can re-install easily, using the ‘root’ account will allow you to run the greatest number of tests.

However, this should not be used to test machines that are in production. A Fuego test can run all kinds of commands, and you should not trust that tests will not destroy your machine (either accidentally or via some malicious intent).

If you don't use 'root', then you can either use an existing account, or create a new account. In most circumstances it is worthwhile to create a new account dedicated to testing. However, you may not have sufficient privileges on your board to do this.

In any event, at this point, decide which account you will use for testing with Fuego, and note it to include in the board file, described later.

6.1.3 3 - Create test directory on target

First, log in to your target board, and create a directory where tests can be run. Usually, you do this as root, and a commonly used directory for this is "/home/fuego". To do this, do the following:

For target with network connectivity :

```
$ ssh root@your_target
<target>$ mkdir /home/fuego
<target>$ exit
```

For target with Serial connectivity :

Use minicom or any other serial terminal tool. Login to the target by giving username and password. Create the directory 'fuego' as below:

```
<target>$ mkdir /home/fuego
```

6.1.4 Create board file

Now, create your board file. The board files reside in <fuego-source-dir>/fuego-ro/boards, and each file has a filename with the name of the board, with the extension ".board".

The easiest way to create a board file is to copy an existing one, and edit the variables to match those of your board. The following instructions are for a board called 'myboard', that has networking access, an ssh daemon running on target, and the ARM architecture.

Do the following:

```
$ cd fuego-ro/boards
$ cp template-dev.board myboard.board
$ vi myboard.board
```

Note: You can use your own editor in place of 'vi'*

Set board parameters

A board file has parameters which define how Fuego interacts with your board. There are lots of different parameters, but the most important to get started quickly (and the most commonly edited) are:

TRANSPORT parameters

Each board needs to specify how Fuego will communicate with it. This is done by specifying a TRANSPORT, and a few variables associated with that transport type.

- TRANSPORT - this specifies the transport to use with the target
 - there are three transport types currently supported: 'ssh', 'serial', 'ttc'
 - Most boards will use the 'ssh' or 'serial' transport type
 - ex: TRANSPORT="ssh"

Most targets require the following:

- LOGIN - specifies the user account to use for Fuego operations
- PASSWORD - specifies the password for that account (if any)

There are some parameters that are specific to individual transports.

For targets using ssh:

- IPADDR
- SSH_PORT
- SSH_KEY

IPADDR is the network address of your board. SSH_PORT is the port where the ssh daemon is listening for connections. By default this is 22, but you should set this to whatever your target board uses. SSH_KEY is the absolute path where an SSH key file may be found (to allow password-less access to a target machine).

An example would be:

- SSH_KEY="/fuego-ro/boards/myboard_id_rsa"

SSH_PORT and SSH_KEY are optional.

For targets using serial:

- SERIAL
- BAUD
- IO_TIME_SERIAL

SERIAL is serial port name used to access the target from the host. This is the name of the serial device node on the host (or in the container).this is specified without the /dev/ prefix.

Some examples are:

- ttyACM0
- ttyACM1
- ttyUSB0

BAUD is the baud-rate used for the serial communication, for eg. “115200”.

IO_TIME_SERIAL is the time required to catch the command’s response from the target. This is specified as a decimal fraction of a second, and is usually very short. A time that usually works is “0.1” seconds.

- ex: IO_TIME_SERIAL=”0.1”

This value directly impacts the speed of operations over the serial port, so it should be adjusted with caution. However, if you find that some operations are not working over the serial port, try increasing this value (in small increments - 0.15, 0.2, etc.)

Note: In the case of `TRANSPORT=”serial”`, Please make sure that docker container and Fuego have sufficient permissions to access the specified serial port. You may need to modify `docker-create-usb-privileged-container.sh` prior to making your docker image, in order to make sure the container can access the ports.

Also, if check that the host filesystem permissions on the device node (e.g `/dev/ttyACM0` allows access. From inside the container you can try using the `sersh` or `sercp` commands directly, to test access to the target.

For targets using `ttc`:

- `TTC_TARGET`

`TTC_TARGET` is the name of the target used with the ‘`ttc`’ command.

Other parameters

- `BOARD_TESTDIR`
- `ARCHITECTURE`
- `TOOLCHAIN`
- `DISTRIB`
- `BOARD_CONTROL`

The `BOARD_TESTDIR` directory is an absolute path in the filesystem on the target board where the Fuego tests are run. Normally this is set to something like “`/home/fuego`”, but you can set it to anything. The user you specify for `LOGIN` should have access rights to this directory.

The `ARCHITECTURE` is a string describing the architecture used by toolchains to build the tests for the target.

The `TOOLCHAIN` variable indicates the toolchain to use to build the tests for the target. If you are using an ARM target, set this to “`debian-armhf`”. This is a default ARM toolchain installed in the docker container, and should work for most ARM boards.

If you are not using ARM, or for some reason the pre-installed arm toolchains don’t work for the Linux distribution installed on your board, then you will need to install your own SDK or toolchain. In this case, follow the steps in [Adding a toolchain](#), then come back to this step and set the `TOOLCHAIN` variable to the name you used for that operation.

For other variables in the board file, see the section below.

The `DISTRIB` variable specifies attributes of the Linux distribution running on the board, that are used by Fuego. Currently, this is mainly used to tell Fuego what kind of system logger the operating system on the board has. Here are some options that are available:

- `base.dist` - a “standard” distribution that implements `syslogd`-style system logging. It should have the commands: `logread`, `logger`, and `/var/log/messages`
- `nologread.dist` - a distribution that has no ‘`logread`’ command, but does have `/var/log/messages`

- `nosyslogd.dist` - a distribution that does not have `syslogd`-style system logging.

If `DISTRIB` is not specified, Fuego will default to using “`nosyslogd.dist`”.

The `BOARD_CONTROL` variable specifies the name of the system used to control board hardware operations. When Fuego is used in conjunction with board control hardware, it can automate more testing functionality. Specifically, it can reboot the board, or re-provision the board, as needed for testing. As of the 1.3 release, Fuego only supports the ‘`ttc`’ board control system. Other board control systems will be introduced and supported over time.

6.1.5 Add node to Jenkins interface

Finally, add the board in the Jenkins interface.

In the Jenkins interface, boards are referred to as “Nodes”.

You can see a list of the boards that Fuego knows about using:

- `$ ftc list-boards`

When you run this command, you should see the name of the board you just created.

You can see the nodes that have already been installed in Jenkins with:

- `$ ftc list-nodes`

To actually add the board as a node in jenkins, inside the docker container, run the following command at a shell prompt:

- `$ ftc add-nodes -b <board_name>`

6.2 Board-specific test variables

The following other variables can also be defined in the board file:

- `MAX_REBOOT_RETRIES`
- `FUEGO_TARGET_TMP`
- `FUEGO_BUILD_FLAGS`

See [Variables](#) for the definition and usage of these variables.

6.2.1 General Variables

6.2.2 File System test variables (SATA, USB, MMC)

If running filesystem tests, you will want to declare the Linux device name and mountpoint path, for the filesystems to be tested. There are three different device/mountpoint options available depending on the testplan you select (SATA, USB, or MMC). Your board may have all of these types of storage available, or only one.

To prepare to run a test on a filesystem on a sata device, define the SATA device and mountpoint variables for your board.

For example, if you had a SATA device with a mountable filesystem accessible on device `/dev/sdb1`, and you have a directory on your target of `/mnt/sata` that can be used to mount this device at, you could declare the following variables in your board file.

- `SATA_DEV="/dev/sdb1"`
- `SATA_MP="/mnt/sata"`

You can define variables with similar names (USB_DEV and USB_MP, or MMC_DEV and MMC_MP) for USB-based filesystems or MMC-based filesystems.

6.2.3 LTP test variables

LTP (the Linux Test Project) test suite is a large collection of tests that require some specialized handling, due to the complexity and diversity of the suite. LTP has a large number of tests, some of which may not work correctly on your board. Some of the LTP tests depend on the kernel configuration or on aspects of your Linux distribution or your configuration.

You can control whether the LTP posix test succeeds by indicating the number of positive and negative results you expect for your board. These numbers are indicated in test variables in the board file:

- LTP_OPEN_POSIX_SUBTEST_COUNT_POS
- LTP_OPEN_POSIX_SUBTEST_COUNT_NEG

You should run the LTP test yourself once, to see what your baseline values should be, then set these to the correct values for your board (configuration and setup).

Then, Fuego will report any deviation from your accepted numbers, for LTP tests on your board.

LTP may also use these other test variables defined in the board file:

- FUNCTIONAL_LTP_HOMEDIR - If this variable is set, it indicates where a pre-installed version of LTP resides in the board's filesystem. This can be used to avoid a lengthy deploy phase on each execution of LTP.
- FUNCTIONAL_LTP_BOARD_SKIPLIST - This variable has a list of individual LTP test programs to skip.

ADDING A TOOLCHAIN

7.1 Introduction

In order to build tests for your target board, you need to install a toolchain (often in the form of an SDK) into the Fuego system, and let Fuego know how to access it.

Adding a toolchain to Fuego consists of these steps:

- 1. obtain (generate or retrieve) the toolchain
- 2. copy the toolchain to the container
- 3. install the toolchain inside the container
- 4. create a `-tools.sh` file for the toolchain
- 5. reference the toolchain in the appropriate board file

7.2 Obtain a toolchain

First, you need to obtain a toolchain that will work with your board. You should have a toolchain that produces software which will work with the Linux distribution on your board. This is usually obtained from your build tool, if you are building the distribution yourself, or from your semiconductor supplier or embedded Linux OS vendor, if you have been provided the Linux distribution from an external source.

7.2.1 Installing a Debian cross-toolchain target

If you are using an Debian-based target, then to get started, you may use a script to install a cross-compiler toolchain into the container. For example, for an ARM target, you might want to install the Debian `armv7hf` toolchain. You can even try a Debian toolchain with other Linux distributions. However, if you are not using Debian on your target board, there is no guarantee that this will produce correct software for your board. It is much better to install your own SDK for your board into the fuego system.

To install a Debian cross toolchain into the container, get to the shell prompt in the container and use the following script:

- `/fuego-ro/toolchains/install_cross_toolchain.sh`

To use the script, pass it the argument naming the cross-compile architecture you are using. Available values are:

- `arm64 armel armhf mips mipsel powerpc ppc64el`

Execute the script, inside the docker container, with a single command-line option to indicate the cross-toolchain to install. You can use the script more than once, if you wish to install multiple toolchains.

Example:

- `# /fuego-ro/toolchains/install_cross_toolchain.sh armhf`

The Debian packages for the specified toolchain will be installed into the docker container.

7.2.2 Building a Yocto Project SDK

When you build an image in the Yocto Project, you can also build an SDK to go with that image using the ‘`-c do_populate_sdk`’ build step with bitbake.

To build the SDK in Yocto Project, inside your yocto build directory do:

- `bitbake <image-name> -c do_populate_sdk`

This will build an SDK archive (containing the toolchain, header files and libraries needed for creating software on your target, and put it into the directory `<build-root>/tmp/deploy/sdk/`

For example, if you are building the ‘core-image-minimal’ image, you would execute:

```
$ bitbake core-image-minimal -c do_populate_sdk
```

At this step look in `tmp/deploy/sdk` and note the name of the sdk install package (the file ending with `.sh`).

7.3 Install the SDK in the docker container

To allow fuego to use the SDK, you need to install it into the fuego docker container. First, transfer the SDK into the container using `docker cp`.

With the container running, on the host machine do:

- `docker ps` (note the container id)
- `docker cp tmp/deploy/sdk/<sdk-install-package> <container-id>:/tmp`

This last command will place the SDK install package into the `/tmp` directory in the container.

Now, install the SDK into the container, wherever you would like. Many toolchains install themselves under `/opt`.

At the shell inside the container, run the SDK install script (which is a self-extracting archive):

- `/tmp/poky-...sh`
 - during the installation, select a toolchain installation location, like: `/opt/poky/2.0.1`

These instructions are for an SDK built by the Yocto Project. Similar instructions would apply for installing a different toolchain or SDK. That is, get the SDK into the container, then install it inside the container.

7.4 Create a -tools.sh file for the toolchain

Now, fuego needs to be told how to interact with the toolchain. During test execution, the fuego system determines what toolchain to use based on the value of the TOOLCHAIN variable in the board file for the target under test. The TOOLCHAIN variable is a string that is used to select the appropriate ‘<TOOLCHAIN>-tools.sh’ file in /fuego-ro/toolchains.

You need to determine a name for this TOOLCHAIN, and then create a file with that name, called \$TOOLCHAIN-tools.sh. So, for example if you created an SDK with poky for the qemuarm image, you might call the TOOLCHAIN “poky-qemuarm”. You would create a file called “poky-qemuarm-tools.sh”

The -tools.sh file is used by Fuego to define the environment variables needed to interact with the SDK. This includes things like CC, AR, and LD. The complete list of variables that this script needs to provide are described on the page [[tools.sh]]

Inside the -tools.sh file, you execute instructions that will set the environment variables needed to build software with that SDK. For an SDK built by the Yocto Project, this involves setting a few variables, and calling the environment-setup... script that comes with the SDK. For SDKs from other sources, you can define the needed variables by directly exporting them.

Here is an example of the tools.sh script for poky-qemuarm. This is in the sample file /fuego-ro/toolchains/poky-qemuarm-tools.sh:

```
# fuego toolchain script
# this sets up the environment needed for fuego to use a
# toolchain
# this includes the following variables:
# CC, CXX, CPP, CXXCPP, CONFIGURE_FLAGS, AS, LD, ARCH
# CROSS_COMPILE, PREFIX, HOST, SDKROOT
# CFLAGS and LDFLAGS are optional
#
# this script is sourced by /fuego-ro/toolchains/tools.sh

POKY_SDK_ROOT=/opt/poky/2.0.1
export SDKROOT=${POKY_SDK_ROOT}/sysroots/
armv5e-poky-linux-gnueabi

# the Yocto project environment setup script changes PATH so
# that python uses
# libs from sysroot, which is not what we want, so save the
# original path
# and use it later
ORIG_PATH=$PATH

PREFIX=arm-poky-linux-gnueabi
source ${POKY_SDK_ROOT}/environment-setup-armv5e-
poky-linux-gnueabi

HOST=arm-poky-linux-gnueabi

# don't use PYTHONHOME from environment setup script
unset PYTHONHOME
env -u PYTHONHOME
```

7.5 Reference the toolchain in a board file

Now, to use that SDK for building test software for a particular target board, set the value of the `TOOLCHAIN` variable in the board file for that target.

Edit the board file:

- `vi /fuego-ro/boards/myboard.board`

And add (or edit) the line:

- `TOOLCHAIN="poky-qemuarm"`

7.6 Notes

7.6.1 Python execution

You may notice that some of the example scripts set the environment variable `ORIG_PATH`. This is used by the function `[[function_run_python|run_python]]` internally to execute the container's default python interpreter, instead of the interpreter that was built by the Yocto Project.

ADDING TEST JOBS TO JENKINS

Before performing any tests with Fuego, you first need to add Jenkins jobs for those tests in Jenkins.

To add jobs to Jenkins, you use the ‘ftc’ command line tool.

Fuego comes with over a hundred different tests, and not all of them will be useful for your environment or testing needs.

In order to add jobs to Jenkins, you first need to have created a Jenkins node for the board for which you wish to add the test. If you have not already added a board definition, or added your board to Jenkins, please see: [Adding a board](#)

Once your board is defined as a Jenkins node, you can add test jobs for it.

There are two ways of adding test jobs, individually, and using testplans. In both cases, you use the ‘ftc add-jobs’ command.

8.1 Selecting tests or plans

The list of all tests that are available can be seen by running the command ‘ftc list-tests’.

Run this command inside the docker container, by going to the shell prompt inside the Fuego docker container, and typing

```
(container_prompt)$ ftc list-tests
```

To see the list of plans that come pre-configured with Fuego, use the command ‘ftc list-plans’.

```
(container_prompt)$ ftc list-plans
```

A plan lists a set of tests to execute. You can examine the list of tests that a testplan includes, by examining the testplan file. The testplan files are in JSON format, and are in the directory `fuego-core/overlays/testplans`.

8.2 Adding individual tests

To add an individual test, add it using the ‘ftc add-jobs’ command. For example, to add the test “Functional.hello_world” for the board “beaglebone”, you would use the following command:

```
(container prompt)$ ftc add-job -b beaglebone -t  
Functional.hello_world
```

8.2.1 Configuring job options

When Fuego executes a test job, several options are available to control aspects of job execution. These can be configured on the ‘ftc add-job’ command line.

The options available are:

- timeout
- rebuild flag
- reboot flag
- precleanup flag
- postcleanup flag

See ‘ftc add-jobs help’ for details about these options and how to specify them.

8.2.2 Adding tests for more than one board

If you want to add tests for more than one board at a time, you can do so by specifying multiple board names after the ‘-b’ option with ‘ftc add-jobs’. Board names should be a single string argument, with individual board names separated by commas.

For example, the following would add a job for `Functional.hello_world` to each of the boards `rpi1`, `rpi2` and `beaglebone`.

```
(container prompt)$ ftc add-job -b rpi1,rpi2,beaglebone -t
Functional.hello_world
```

8.3 Adding jobs based on testplans

A testplan is a list of Fuego tests with some options for each one. You can see the list of testplans in your system with the following command:

```
(container prompt)$ ftc list-plans
```

To create a set of jobs related to docker image testing, for the ‘docker’ board on the system, do the following:

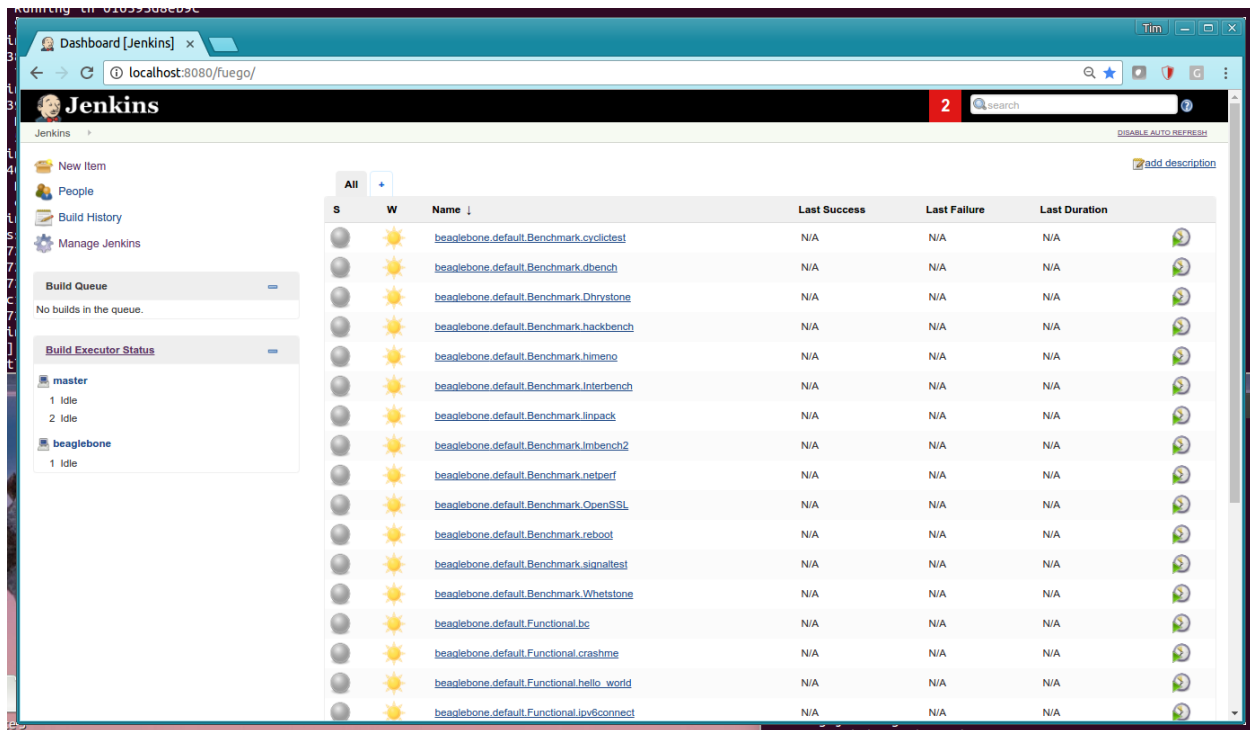
```
(container prompt)$ ftc add-jobs -b docker -p testplan_docker
```

To create a set of jobs for a board called ‘beaglebone’, do the following:

```
(container prompt)$ ftc add-jobs -b myboard -p testplan_smoketest
```

The “smoketest” testplan has about 20 tests that exercise a variety of features on a Linux system. After running these commands, a set of jobs will appear in the Jenkins interface.

Once this is done, your Jenkins interface should look something like this:



The screenshot shows the Jenkins web interface. The top navigation bar includes 'Dashboard [Jenkins]', a search bar, and a '2' badge. The left sidebar contains links for 'New Item', 'People', 'Build History', and 'Manage Jenkins'. Below these are sections for 'Build Queue' (showing 'No builds in the queue') and 'Build Executor Status' (showing 'master' with 1 idle and 2 idle executors, and 'beaglebone' with 1 idle executor).

The main content area displays a table of jobs. The table has columns for 'S' (Success), 'W' (Warning), 'Name', 'Last Success', 'Last Failure', and 'Last Duration'. Each row represents a job, with a status icon (sun for success, cloud for warning) and a link to the job's details.

S	W	Name	Last Success	Last Failure	Last Duration
		beaglebone.default.Benchmark.cyclistest	N/A	N/A	N/A
		beaglebone.default.Benchmark.dbench	N/A	N/A	N/A
		beaglebone.default.Benchmark.Dhrystone	N/A	N/A	N/A
		beaglebone.default.Benchmark.hackbench	N/A	N/A	N/A
		beaglebone.default.Benchmark.himeno	N/A	N/A	N/A
		beaglebone.default.Benchmark.interbench	N/A	N/A	N/A
		beaglebone.default.Benchmark.linpack	N/A	N/A	N/A
		beaglebone.default.Benchmark.lmbench2	N/A	N/A	N/A
		beaglebone.default.Benchmark.netperf	N/A	N/A	N/A
		beaglebone.default.Benchmark.OpenSSL	N/A	N/A	N/A
		beaglebone.default.Benchmark.reboot	N/A	N/A	N/A
		beaglebone.default.Benchmark.signaltest	N/A	N/A	N/A
		beaglebone.default.Benchmark.Whetstone	N/A	N/A	N/A
		beaglebone.default.Functional.bc	N/A	N/A	N/A
		beaglebone.default.Functional.crashme	N/A	N/A	N/A
		beaglebone.default.Functional.hello_world	N/A	N/A	N/A
		beaglebone.default.Functional.ipv6connect	N/A	N/A	N/A

ADDING VIEWS TO JENKINS

It is useful to organize your Jenkins test jobs into “views”. These appear as tabs in the main Jenkins interface. Jenkins always provides a tab that lists all of the installed jobs, call “All”. Other views that you create will appear on tabs next to this, on the main Jenkins page.

You can define new Jenkins views using the Jenkins interface, but Fuego provides a command that allows you to easily create views for boards, or for sets of related tests (by name and wildcard), from the Linux command line (inside the container).

The usage line for this command is:

```
Usage: ftc add-view <view-name> [<job_spec>]
```

The view-name parameter indicates the name of the view in Jenkins, and the job-spec parameter is used to select the jobs which appear in that view.

If the job_spec is provided and starts with an ‘=’, then it is interpreted as one or more specific job names. Otherwise, the view is created using a regular expression statement that Jenkins uses to select the jobs to include in the view.

9.1 Adding a board view

By convention, most Fuego users populate their Jenkins interface with a view for each board in their system (well, for labs with a small number of boards, anyway).

The simplest way to add a view for a board is to just specify the board name, like so:

```
(container_prompt)$ ftc add-view myboard
```

When no job specification is provided, the ‘add-view’ command will create one by prefixing the view name with wildcards. For the example above, the job spec would consist of the regular expression “.*myboard.*”.

9.1.1 Customizing regular expressions

Note that if your board name is not unique enough, or is a string contained in some tests, then you might see some test jobs listed that were not specific to that board. For example, if you had a board name “Bench”, then a view you created with the view-name of “Bench”, would also include Benchmarks. You can work around this by specifying a more details regular expression for your job spec.

For example:

```
(container_prompt)$ ftc add-view Bench "Bench.*"
```

This would only include the jobs that started with “Bench” in the “Bench” view. Benchmark jobs for other boards would not be included, since they only have “Benchmark” somewhere in the middle of their job name - not at the beginning.

9.2 Add view by test name regular expression

This command would create a view to show LTP results for multiple boards:

```
(container_prompt)$ ftc add-view LTP
```

This example creates a view for “fuego” tests. This view would include any job that has the word “fuego” as part of it. By convention, all Fuego self-tests have part of their name prefixed with “*fuego_*”.

```
(container_prompt)$ ftc add-view fuego ".*fuego_.*"
```

And the following command will show all the batch jobs defined in the system:

```
(container_prompt)$ ftc add-view *.batch
```

9.3 Add specific jobs

If the job specification starts with “=”, it is a comma-separated list of job names. The job names must be complete, including the board name, spec name and full test name.

```
(container_prompt)$ ftc add-view network-tests =docker.default.  
Functional.ipv6connect,docker.default.Functional.netperf
```

In this command, the view would be named “network-tests”, and it would show the jobs “docker.default.Functional.ipv6connect” and “docker.default.Functional.netperf”.

TEST VARIABLES

10.1 Introduction

When Fuego executes a test, shell environment variables are used to provide information about the test environment, test execution parameters, communications methods and parameters, and other items.

These pieces of information are originate from numerous different places. An initial set of test variables comes in the shell environment from either Jenkins or from the shell in which `ftc` is executed (depending on which one is used to invoke the test).

The information about the board being tested comes primarily from two sources:

- The board file
- The stored board variables file

Additional information comes from the testplan and test spec that are used for this particular test run. Finally, test variables can be defined on the `ftc` command line. These test variables (know as *dynamic variables*, override variables that come from other sources.

Test variables can be simple strings, or they may be shell functions.

When a test is run, Fuego gathers information from all these sources, and makes them available to the test (and uses them itself) to control test execution.

10.2 Board file

The board file contains static information about a board. It is processed by the overlay system, and the values inside it appear as variables in the environment of a test, during test execution.

The board file resides in `/fuego-ro/boards` and the filename ends in the string `".board"`:

- `/fuego-ro/boards/{board_name}.board`

There are a number of variables which are used by the Fuego system itself, and there may also be variables that are used by individual tests.

10.2.1 Common board variables

Here is a list of the foo bar variables which might be found in a board file:

- ARCHITECTURE - specifies the architecture of the board
- BAUD - baud rate for serial device (if using ‘serial’ transport)
- BOARD_TESTDIR - directory on board where tests are executed
- BOARD_CONTROL - the mechanism used to control board hardware (e.g. hardware reboot)
- DISTRIB - filename of distribution overlay file (if not the default)
- IO_TIME_SERIAL - serial port delay parameter (if using ‘serial’ transport)
- IPADDR - network address of the board
- LOGIN - specifies the user account to use for Fuego operations
- PASSWORD - specifies the password for the user account on the board used by Fuego
- PLATFORM - specifies the toolchain to use for the platform
- SATA_DEV - specifies a filesystem device node (on the board) for SATA filesystem tests
- SATA_MP - specifies a filesystem mount point (on the board) for SATA filesystem tests
- SERIAL - serial device on host for board’s serial console (if using ‘serial’ transport)
- SRV_IP - network address of server endpoint, for networking tests (if not the same as the host)
- SSH_KEY - the absolute path to key file with ssh key for password-less ssh operations (e.g. “/fuego-ro/board/myboard_id_rsa”)
- SSH_PORT - network port of ssh daemon on board (if using ssh transport)
- TRANSPORT - this specifies the transport to use with the target
- USB_DEV - specifies a filesystem device node (on the board) for USB filesystem tests
- USB_MP - specifies a filesystem mount point (on the board) for USB filesystem tests

See [Adding a board](#) for more details about these variables.

A board may also have additional variables, including variables that are used for results evaluation for specific tests.

10.3 Overlay system

The overlay system gathers variables from several places, and puts them all together into a single file which is then source’ed into the running test’s environment.

It takes information from:

- The board files (both static and dynamic)
- The testplan
- The test spec
- The overlay files

and combines them all, using a set of priorities, into a single file called `prolog.sh`, which is then source’ed into the running shell environment of the Fuego test being executed.

The overlay system is described in greater detail here: [Overlay_Generation](#)

10.4 Stored variables

Stored board variables are test variables that are defined on a per-board basis, and can be modified and managed under program control.

Stored variables allow the Fuego system, a test, or a user to store information that can be used by tests. This essentially creates an information cache about the board, that can be both manually and programmatically generated and managed.

The information that needs to be held for a particular board depends on the tests that are installed in the system. Thus the system needs to support ad-hoc collections of variables. Just putting everything into the static board file would not scale, as the number of tests increases.

Note: The LAVA test framework has a similar concept called a board dictionary.

One use case for this is to have a “board setup” test, that scans for lots of different items, and populates the stored variables with values that are used by other tests. Some items that are useful to know about a board take time to discover (using e.g. `find` on the target board), and using a board dynamic variable can help reduce the time required to check these items.

The board stored variables are kept in the file:

- `/fuego-rw/boards/{board_name}.vars`

These variables are included in the test by the overlay generator.

10.4.1 Commands for interacting with stored variables

A user or a test can manipulate a board stored variable using the `ftc` command. The following commands can be used to set, query and delete variables:

- `tc query-board` - to see test variables (both regular board variables and stored variables)
- `ftc set-var` - to add or update a stored variable
- `ftc delete-var` - to delete a stored variable

ftc query-board

`ftc query-board` is used to view the variables associated with a Fuego board. You can use the command to see all the variables, or just a single variable.

Note that `ftc query-board` shows the variables for a test that come from both the board file and board stored variables file (that is, both ‘static’ board variables and stored variables). It does not show variables which come from testplans or spec files, as those are specific to a test.

The usage is:

- `ftc query-board <board> [-n <VARIABLE>]`

Examples:

```
$ ftc query-board myboard
$ ftc query-board myboard -n PROGRAM_BC
```

The first example would show all board variables, including functions. The second example would show only the variable `PROGRAM_BC`, if it existed, for board ‘myboard’.

ftc set-var

`ftc set-var` allows setting or updating the value of a board stored variable.

The usage is:

- `ftc set-var <board> <VARIABLE>=<value>`

By convention, variable names are all uppercase, and function names are lowercase, with words separated by underscores.

Example:

```
$ ftc set-var PROGRAM_BC=/usr/bin/bc
```

ftc delete-var

`ftc delete-var` removes a variable from the stored variables file.

Example:

```
$ ftc delete-var PROGRAM_BC
```

10.4.2 Example usage

The test `Functional.fuego_board_check` could detect the path for the `foo` binary, (e.g. `is_on_target foo PROGRAM_FOO`) and call `ftc set-var $NODE_NAME PROGRAM_FOO=$PROGRAM_FOO`. This would stay persistently defined as a test variable, so other tests could use `$PROGRAM_FOO` (with `assert_define`, or in `report` or `cmd` function calls.)

10.4.3 Example Stored variables

Here are some examples of variables that can be kept as stored variables, rather than static variables from the board file:

- `SATA_DEV` = Linux device node for SATA file system tests
- `SATA_MP` = Linux mount point for SATA file system tests
- `LTP_OPEN_POSIX_SUBTEST_COUNT_POS` = expected number of pass results for LTP OpenPosix test
- `LTP_OPEN_POSIX_SUBTEST_COUNT_NEG` = expected number of fail results for LTP OpenPosix test
- `PROGRAM_BC` = path to 'bc' program on the target board
- `MAX_REBOOT_RETRIES` = number of retries to use when rebooting a board

10.5 Spec variables

A test spec can define one or more variables to be used with a test. These are commonly used to control test variations, and are specified in a `spec.json` file.

When a spec file defines a variable associated with a named test spec, the variable is read by the overlay generator on test execution, and the variable name is prefixed with the name of the test, and converted to all upper case.

For example, support a test called `Functional.foo` had a test spec that defined the variable ‘args’ with a line like the following in its `spec.json` file:

```
"default": {
  "args": "-v -p2"
}
```

When the test was run with this spec (the “default” spec), then the variable `FUNCTIONAL_FOO_ARGS` would be defined, with the value “-v -p2”.

See `Test_Specs_and_Plans` for more information about specs and plans.

10.6 Dynamic variables

Another category of variables used during testing are dynamic variables. These variables are defined on the command line of `ftc run-test` using the `--dynamic-vars` option.

The purpose of these variables is to allow scripted variations when running `ftc run-test`. The scripted variables are processed and presented the same way as spec variables, which is to say that the variable name is prefixed with the test name, and converted to all upper case.

For example, if the following command was issued:

- `ftc run-test -b beaglebone -t Functional.foo --dynamic_vars *ARGS=-p*`

then during test execution the variable `FUNCTIONAL_FOO_ARGS` would be defined with the value “-p”.

See `Dynamic Variables` for more information.

10.7 Variable precedence

Here is the precedence of variable definition for Fuego, during test execution:

(from lowest to highest)

1. environment variable (from Jenkins or shell where ‘ftc run-test’ is invoked)
2. board variable (from `fuego-ro/boards/$BOARD.board` file)
3. stored variable (from `fuego-rw/boards/$BOARD.vars` file)
4. spec variable (from `spec.json` file)
5. dynamic variable (from `ftc` command line)
6. core variable (from Fuego scripts)
7. `fuego_test` variable (from `fuego_test.sh`)

Spec and dynamic variables are prefixed with the test name, and converted to upper case. That tends to keep them in a separate name space from the rest of the test variables.

JENKINS USER INTERFACE

By default, Fuego uses the Jenkins continuous integration system to manage boards, tests, logs, and test results.

The Jenkins user interface is web-based. This page shows several screenshots of different pages in the Jenkins interface.

Through this interface, you can see the status of tests that have run, review the logs for tests, and schedule new tests to run on target boards. You also use this interface to add new boards and new tests to the system.

Note that Jenkins objects are:

- nodes
- jobs
- builds
- views

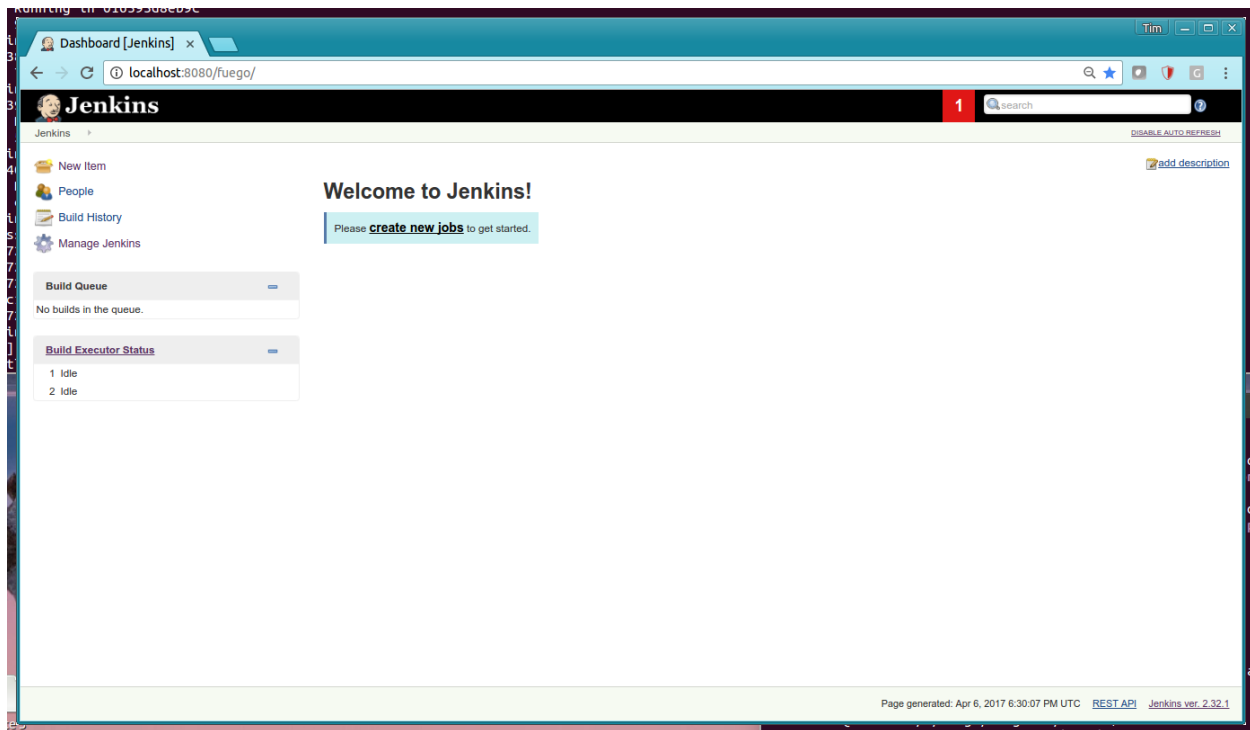
These are different from the Fuego names for the same objects. The first three of these Jenkins objects correspond to the Fuego objects of: **boards**, **tests** and **runs**, respectively.

11.1 Main dashboard

The main dashboard of Jenkins looks like the following:

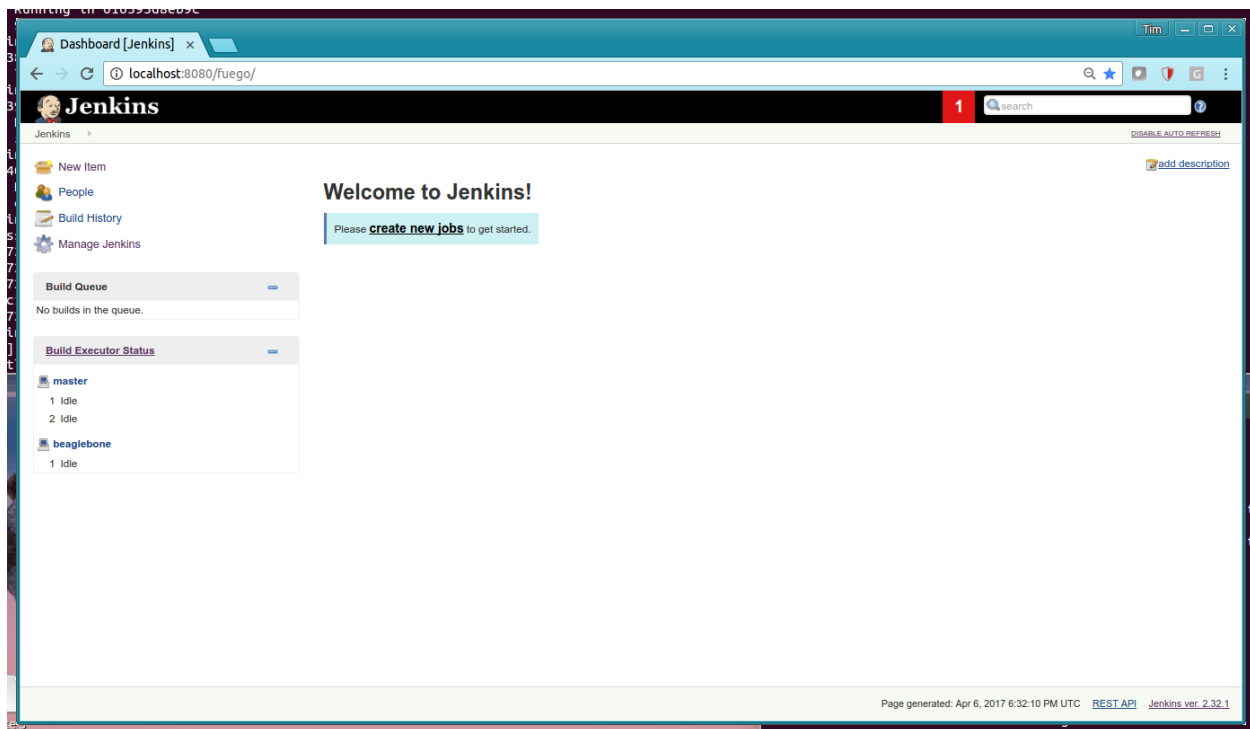
11.1.1 New Installation

When Fuego has just been installed, there is nothing in the Jenkins interface (no nodes, jobs or views). The interface should look something like this:



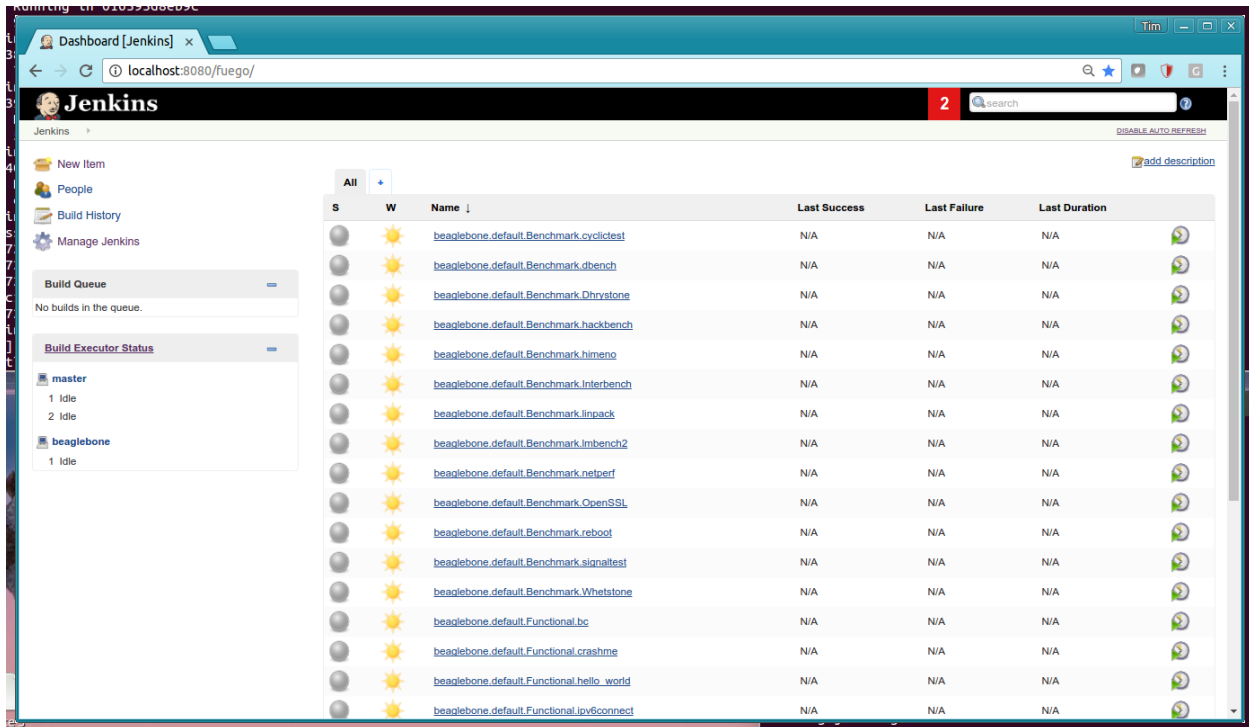
11.1.2 With a single node (board) added

Here is the main dashboard of Jenkins, after a single node (called 'beaglebone' in this case) has been added. Note the node (board) appears in the left sidebar under "Build Executor Status":



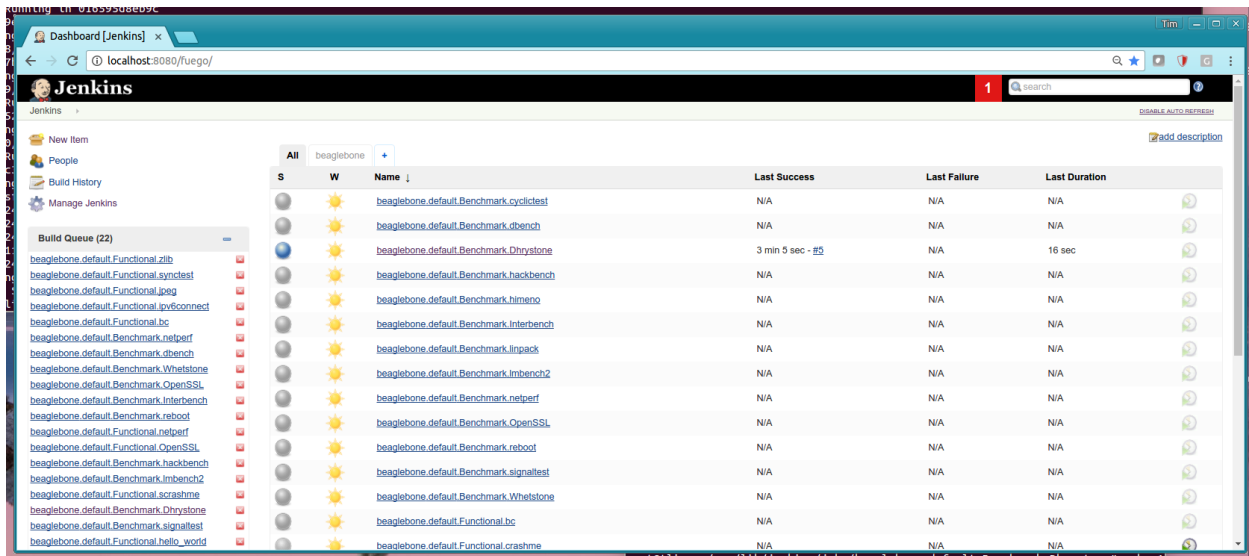
11.1.3 With beaglebone node and jobs

Here is the main dashboard of Jenkins, showing a single node (beaglebone) and jobs for this board.



11.1.4 Dashboard with jobs in Build Queue

Here is the Jenkins dashboard with a lot of jobs in the Build Queue. Note the list of jobs in the left side-bar, in the “Build Queue” pane.



11.2 Node pages

If you click on the node in the **Build Executor Status** pane, then you can see a list of the jobs associated with a node.

11.2.1 Node status page

Here is the status for the beaglebone node.

The screenshot shows the Jenkins web interface for the 'Agent beaglebone'. The main content area is titled 'Agent beaglebone' and lists 'Projects tied to beaglebone'. The table below shows the status of various benchmarks.

S	W	Name	Last Success	Last Failure	Last Duration
●	☀	beaglebone.default.Benchmark.cyclidest	37 min - #2	N/A	19 sec
●	☀	beaglebone.default.Benchmark.dbench	N/A	N/A	N/A
●	☀	beaglebone.default.Benchmark.Dhrystone	37 min - #7	N/A	16 sec
●	☀	beaglebone.default.Benchmark.hackbench	37 min - #2	N/A	12 sec
●	☀	beaglebone.default.Benchmark.himeno	39 min - #2	N/A	1 min 6 sec
●	☀	beaglebone.default.Benchmark.interbench	N/A	N/A	N/A
●	☀	beaglebone.default.Benchmark.linpack	38 min - #2	N/A	29 sec
●	☀	beaglebone.default.Benchmark.lmbench2	N/A	N/A	N/A
●	☀	beaglebone.default.Benchmark.netperf	N/A	N/A	N/A
●	☀	beaglebone.default.Benchmark.OpenSSL	N/A	N/A	N/A
●	☀	beaglebone.default.Benchmark.reboot	N/A	N/A	N/A
●	☀	beaglebone.default.Benchmark.signaltest	37 min - #2	N/A	16 sec
●	☀	beaglebone.default.Benchmark.Whetstone	35 min - #3	N/A	1 min 14 sec

The left sidebar shows the 'Build Executor Status' pane with '1 Idle'.

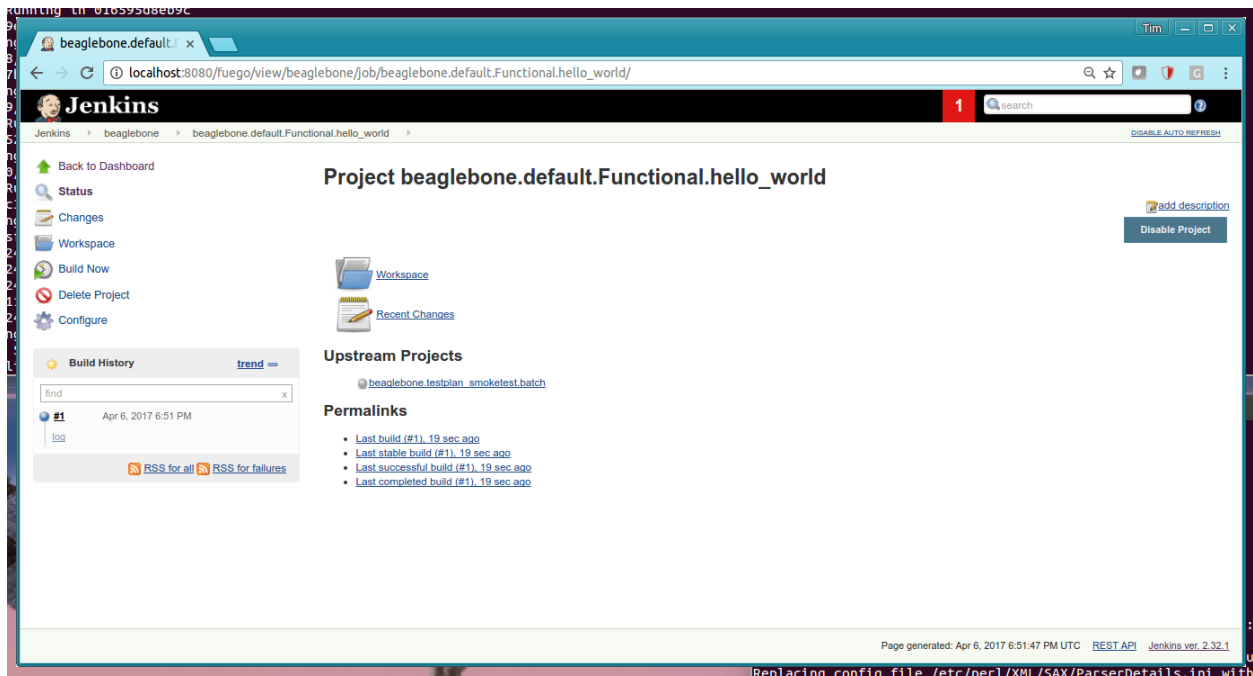
11.3 Job pages

If you click on a job in the Jenkins interface, you can see information about an individual job. This page shows information about the status of the job, including a Build History for the job (in the left sidebar).

You can start a job by clicking on the “Build Now” button in the left menu.

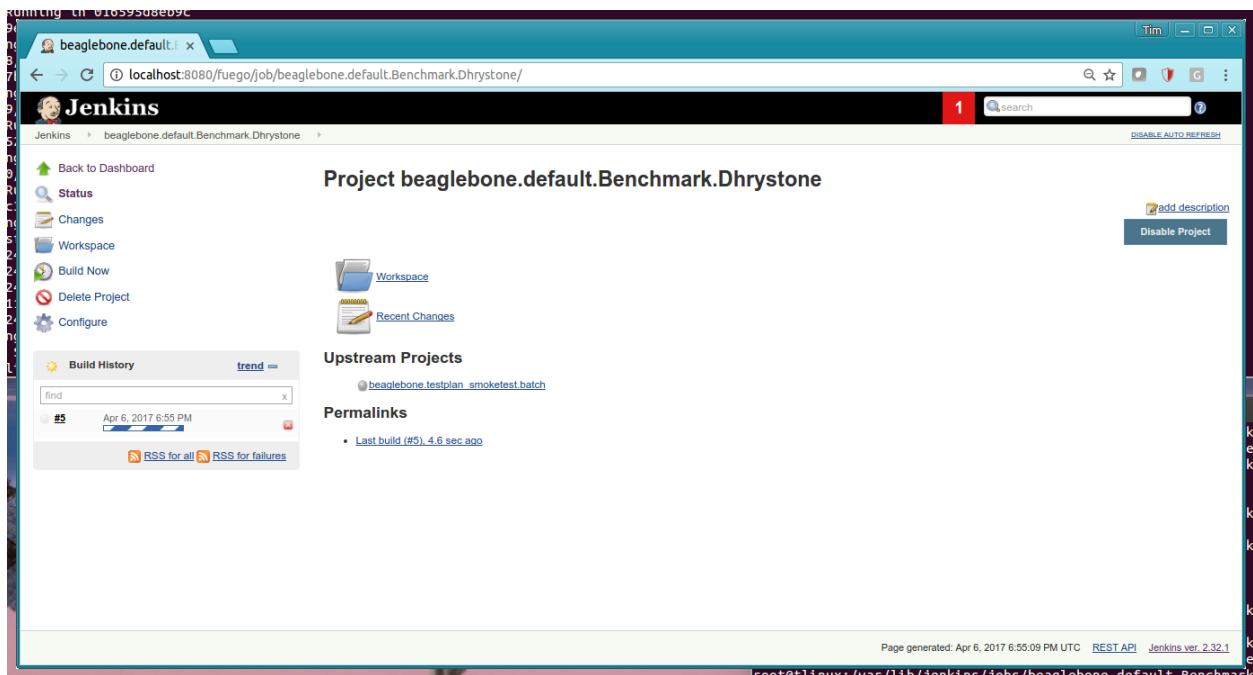
11.3.1 Functional job status page

Here is a page showing the status information for a Functional test called ‘hello_world’. The main area of the screen has information about the last successful and failed builds of the test. Note the left sidebar pane with “Build History”, to see individual test execution results.



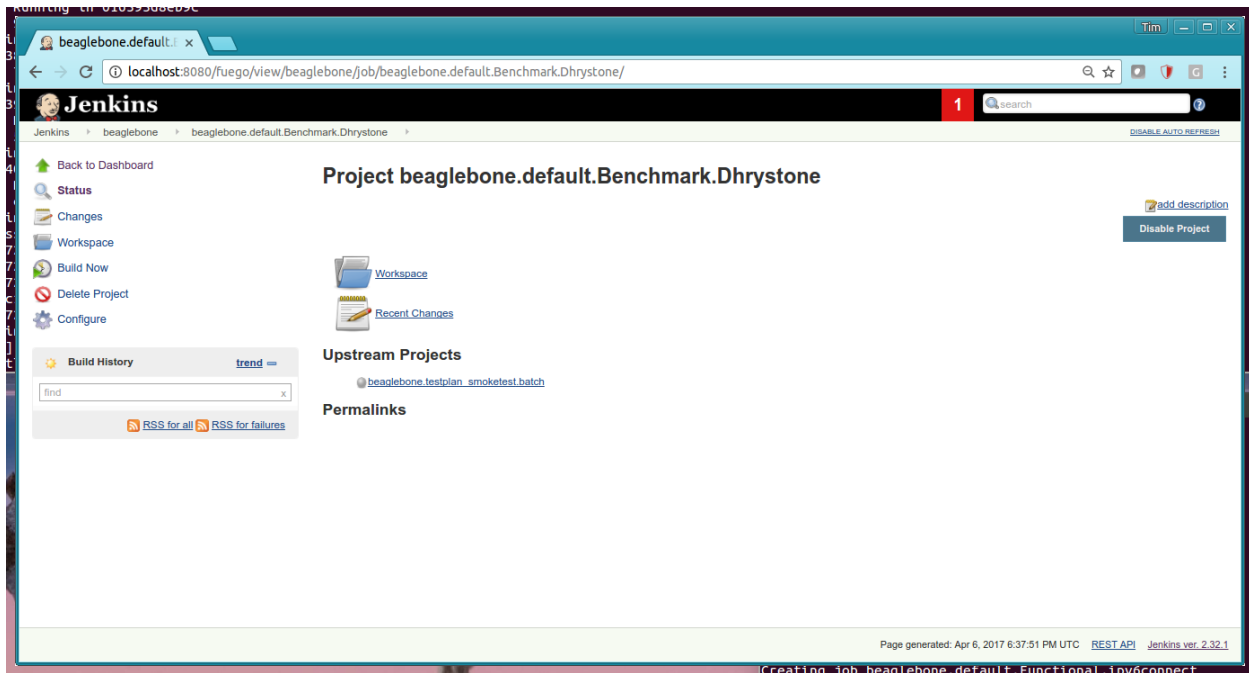
11.3.2 Benchmark job - starting a build

Here is a picture of a job build being started. Note the progress bar in the **Build History** pane in the left sidebar.



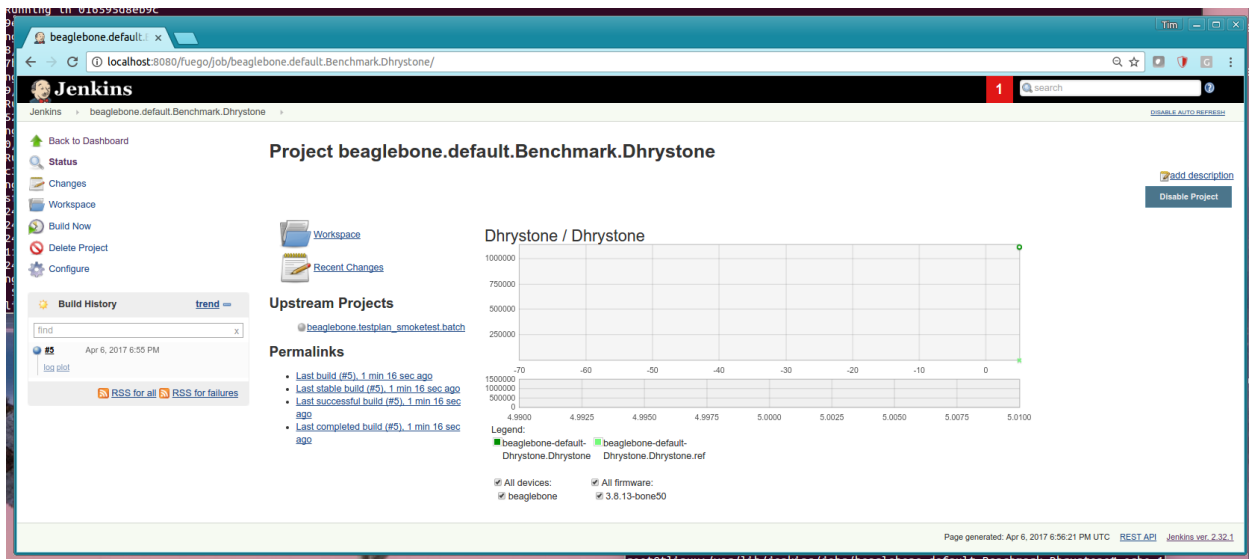
11.3.3 Benchmark job - before successful execution

Before a Benchmark job completes it has no data to plot on it's chart, and appears similar to a Functional Job status page:



11.3.4 Benchmark job - with plot of metrics

Normally, a Benchmark page shows one or more plots showing the values for data returned by this benchmark.

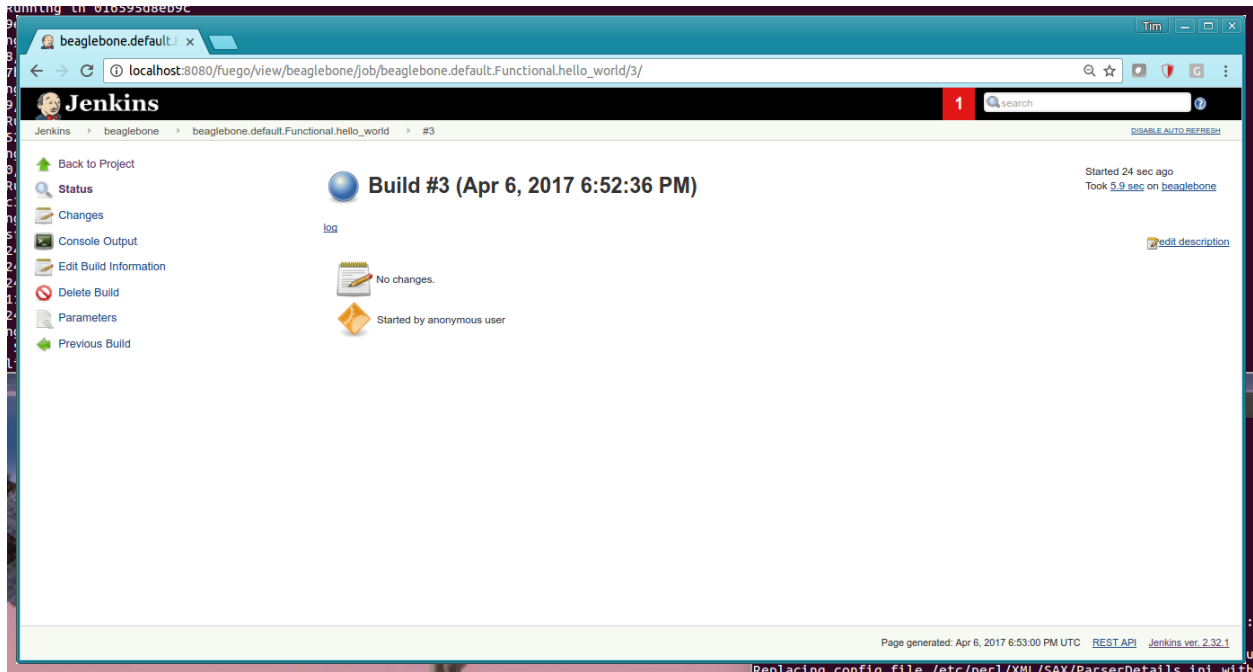


11.4 Build pages

A build page shows the results of a single execution of a job (test) on a board. You can click on the build number in the Jenkins interface to see this page.

11.4.1 Results from a job build

Here are the results from the execution of the “hello world” job. This was the results of running the Fuego test “Functional.hello_world” on a beaglebone board.



11.4.2 Test log results

You can examine the different logs for each test. Each test produces a log from the program that ran on the board. This is available by following a link called “log” from the ‘build’ page for that test run. You can see the console log, which shows the output of commands for this test, by clicking on “console log” in the build interface (or the build drop-down menu in the **Build History** list).

Drhystone test log

Here are results from a run of the Drhystone test on a beaglebone board:

```

Dhrystone Benchmark, Version 2.1 (Language: C)
Program compiled without 'register' attribute
Please give the number of runs through the benchmark:
Execution starts, 10000000 runs through Dhrystone
Execution ends

Final values of the variables used in the benchmark:

Int_Glob:      5
  should be:   5
Bool_Glob:     1
  should be:   1
Ch_1_Glob:     A
  should be:   A
Ch_2_Glob:     B
  should be:   B
Arr_1_Glob[8]: 7
  should be:   7
Arr_2_Glob[8][7]: 10000010
  should be:   Number_Of_Runs + 10
Ptr_Glob->
  Ptr_Comp:    151560
  should be:   (implementation-dependent)
  Discr:       0
  should be:   0
Enum_Comp:     2
  should be:   2
Int_Comp:      17
  should be:   17
Str_Comp:      DHRYSTONE PROGRAM, SOME STRING
  should be:   DHRYSTONE PROGRAM, SOME STRING
Next_Ptr_Glob->
  Ptr_Comp:    151560
  should be:   (implementation-dependent), same as above
  Discr:       0
  should be:   0
Enum_Comp:     1
  should be:   1
Int_Comp:      18
  should be:   18
Str_Comp:      DHRYSTONE PROGRAM, SOME STRING
  should be:   DHRYSTONE PROGRAM, SOME STRING
Tnt_1_Inc:     5
  
```

Jenkins Console log

Here is the console log for a test executed on the beaglebone:

```

Started by upstream project "beaglebone.testplan_smoketest.batch" build number 2
originally caused by:
  Started by user anonymous
Building remotely on beaglebone in workspace /fuego-rw/buildzone
[buildzone] $ /bin/sh -xe /tmp/hudson1854422796647547463.sh
+ export Reboot=false
+ Reboot=false
+ export Rebuild=true
+ Rebuild=true
+ export Target_PreCleanup=true
+ Target_PreCleanup=true
+ export Target_PostCleanup=true
+ Target_PostCleanup=true
+ export TESTDIR=Functional.hello_world
+ TESTDIR=Functional.hello_world
+ export TESTNAME=hello_world
+ TESTNAME=hello_world
+ export TESTSPEC=default
+ TESTSPEC=default
+ timeout --signal=9 100m /bin/bash /fuego-core/engine/tests/Functional.hello_world/hello_world.sh
Using nosyslogdist overlay
#### doing fuego phase: pre_test ####
AS=arm-linux-gnueabihf-as
LDFLAGS=-sysroot / -lm
AR=arm-linux-gnueabihf-ar
CXXCPP=arm-linux-gnueabihf-cpp
FUEGO_RM=/fuego-rw
HOST=arm-linux-gnueabihf
TERM=xterm
SHELL=/bin/bash
HUDSON_SERVER_COOKIE=f41155c50dalf1cd
  
```

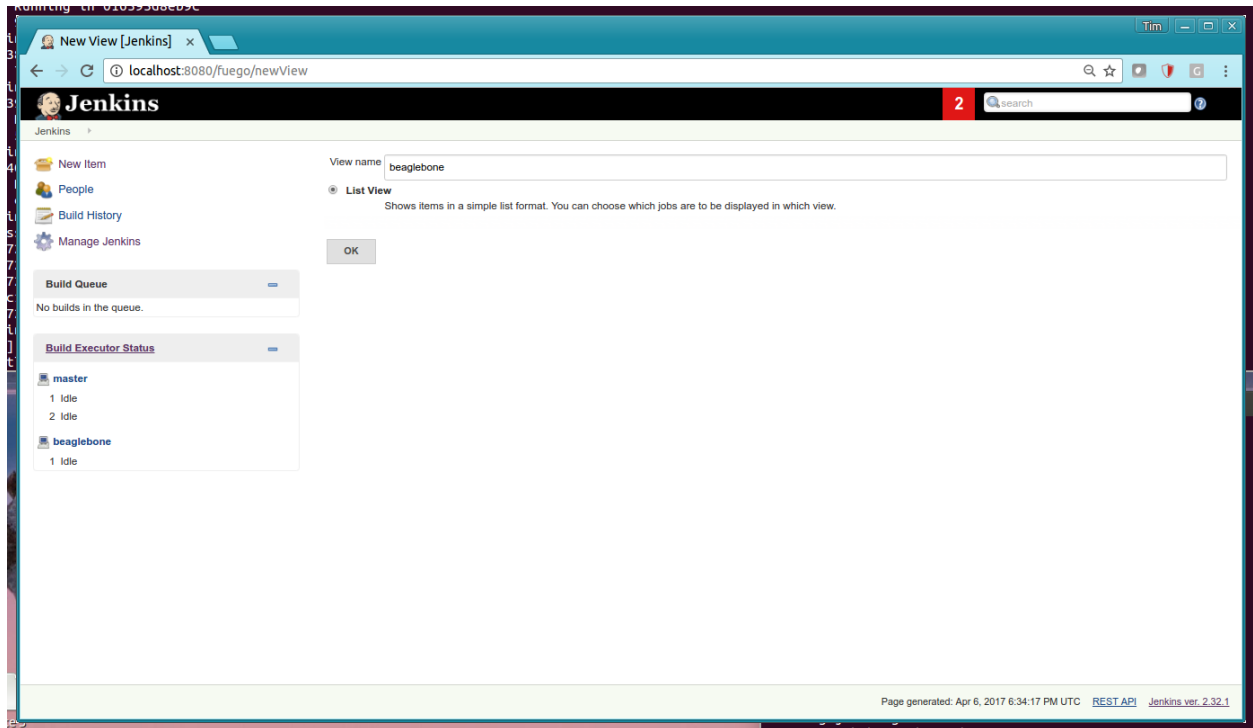
11.5 View pages

A **view** is an area in the main dashboard that shows a collection of jobs with a particular set of status columns for each job. They appear as tabs in the main dashboard view of Jenkins. You can create your own view to see a subset of the jobs that are available in Jenkins

Here are some screen shots showing how to add a new view to Jenkins.

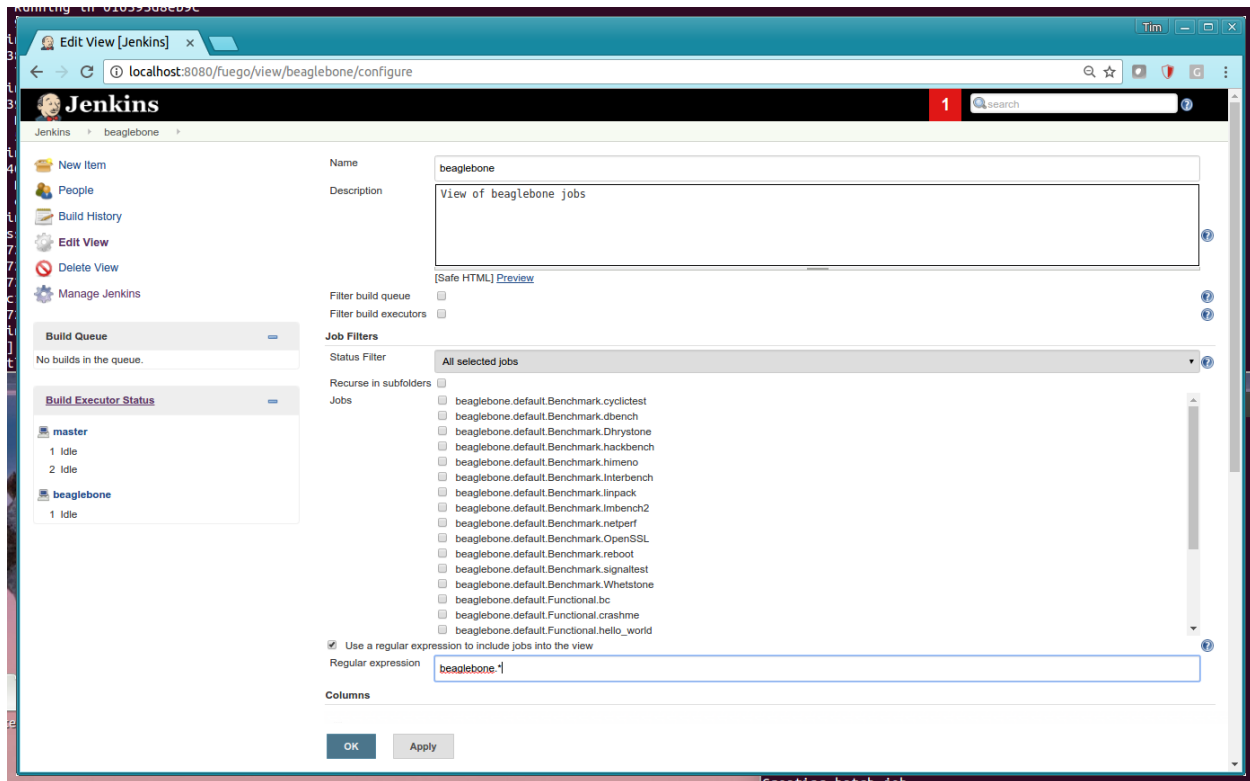
11.5.1 Screen to add a new view

Here is the screen to add a new view to Jenkins.



11.5.2 Screen to configure the view settings

Here is the screen to configure view settings. Note the use of a regular expression to control what jobs to see in this view. You can also control what status columns to display in the view.



11.6 Other Jenkins pages

11.6.1 Build History

The global build history page is available by clicking on the **Build History** link in the main dashboard page of Jenkins. It shows the execution time and status for a recent time period.

Build History on beaglebone

Build	Time Since ↑	Status
beaglebone.default.Benchmark.reboot #4	27 min	aborted
beaglebone.default.Functional.netperf #3	33 min	aborted
beaglebone.default.Benchmark.reboot #3	33 min	aborted
beaglebone.default.Benchmark.dbench #4	33 min	aborted
beaglebone.default.Benchmark.interbench #3	38 min	aborted

11.6.2 Jenkins management

You can manage Jenkins using the **Manage Jenkins** page, available from the top-level dashboard page in Jenkins. From here you can update Jenkins itself, install or remove plugins, and perform other management operations for Jenkins.

Manage Jenkins

- Configure System**: Configure global settings and paths.
- Configure Global Security**: Secure Jenkins; define who is allowed to access/use the system.
- Global Tool Configuration**: Configure tools, their locations and automatic installers.
- Reload Configuration from Disk**: Discard all the loaded data in memory and reload everything from file system. Useful when you modified config files directly on disk.
- Manage Plugins**: Add, remove, disable or enable plugins that can extend the functionality of Jenkins.
- System Information**: Displays various environmental information to assist trouble-shooting.
- System Log**: System log captures output from java.util.logging output related to Jenkins.
- Load Statistics**: Check your resource utilization and see if you need more computers for your builds.
- Jenkins CLI**: Access/manage Jenkins from your shell, or from your script.
- Script Console**: Executes arbitrary script for administration/trouble-shooting/diagnostics.
- Manage Nodes**: Add, remove, control and monitor the various nodes that Jenkins runs jobs on.
- About Jenkins**

COMMAND LINE TOOL - FTC

12.1 Introduction

Fuego comes with a command line tool, called `ftc`, that is used to perform a variety of tasks.

`ftc` has many sub-commands for performing different operations, but they can be grouped into different categories, including the following:

- Essential commands
- Commands for inspecting results
- Commands for managing boards
- Commands for working with Jenkins
- Commands for working with a Fuego (or other) server
- Miscellaneous commands

Note that you can get a list of all the commands supported by `ftc`, by using `ftc help`.

For help on any individual sub-command, use `ftc help <command>`.

12.1.1 General FTC options

Some options to `ftc` commands are used consistently throughout the tool. That is, there are some global options, which apply to any sub-command, such as `-v` for verbose mode. Or, some options are always specified with the same argument letter, such as `-b` to specify the board for the operation.

Here are options that are commonly used with `ftc` commands:

- `-v` = verbose mode - this means to report more information than usual
- `-q` = quiet mode - this means to report less information than usual, sometimes making the command completely silent. Quiet mode is often used to omit header data and make the output suitable for parsing by other tools.
- `-h` = help - show help for the specified sub-command
- `-debug` = run in debug mode - this makes `ftc` produce debug information while it runs
- `-b` = specify the board - used when a command takes a board argument
- `-t` = specify the test - used when a command test a test argument.

Note: For commands that require a board, you can specify the board a few different ways. The most common is to use the ‘-b’ option followed by the board name. You can omit this if you have a board name specified in the environment variable FUEGO_BOARD. Also, you can set a default board name in your fuego configuration, which is found in fuego-ro/conf/fuego.conf.

Some commands (such as `add-board`, `add-job`, or `run-test`) can operate on multiple boards. In this case, use the ‘-b’ option and specify multiple board names separated by commas after the ‘-b’, like this: `ftc add-board -b board1,board2`

In some cases, you don’t need to specify the full board name, but can provide a short string, as long as the string is sufficient to identify the desired board uniquely. For example, if you had only a single board name that started with a ‘q’ (like ‘qemu-arm’), you could use something like this: `ftc add-board -b q`

Note: A test name is specified using either its full name, which includes the test type, or its shortened name, which is the portion of the name after the type. For example, for the test `Functional.hello_world`, the full test name is “Functional.hello_world”, and the short name is just “hello_world”. Either name may be specified with the ‘-t’ option, although it is usually more convenient to use the short name.

In case there is both a Benchmark and Functional test with the same short name, the full name must be used with ‘-t’ to make it clear which test is referred to.

12.2 Commands groups

12.2.1 Essential commands

Here are the commands that are essential for adding information to the Jenkins interface (nodes and jobs), for querying the available boards, tests, and test specs, and for executing a test.

- `add-nodes`
- `add-jobs`
- `list-boards`
- `list-tests`
- `list-specs`
- `run-test`

These core commands are used for setup of jobs in Jenkins, and for finding out (or trying to remember), the test elements that are available on your system for test execution. Use `ftc add-nodes` and `ftc add-jobs` to populate the Jenkins interface with nodes and jobs (respectively). See the section *Jenkins-related commands* below for more information about these commands, or see *Adding a Board* or *Adding test jobs to Jenkins* for more details.

The most important `ftc` command is `ftc run-test`. This is the command used to actually run a test. Details about `ftc run-test` can be found in the section *ftc run-test* below.

12.2.2 Commands for inspecting results

`ftc` provides a few commands for inspecting the results from operations. Usually, Jenkins is used to show visualization of results. However, you can also see what tests have been run, and what the results from those tests were, using the following commands:

- `list-runs`
- `gen-report`

Use the `list-runs` command to see a list of all test runs on the system. The list of runs can be filtered by various criteria.

`FIXTHIS` - reference `--where` clause description section

Use the `gen-report` command to generate reports of test results data.

Usually, test results are examined in the *Jenkins User interface*. However, you can also generate lists of results at the command line using `ftc gen-report`

This command gives you control over the results that are reported, as well as the content (exact fields and headers) and format of the report.

In summary, `ftc gen-report` can:

- select the test runs to report results from
- select the header fields to show in the report
- select the data (result) fields to show in the report
- filter the data by results (for example showing only failures)
- select the format of the report
- specify the output location for the report

See the section *Generating Reports* for details about this command and its options, and overall information about generating reports from test run data.

12.2.3 Board management

These commands have to do with managing boards (defined on the local machine):

- `list-boards`
- `query-board`
- `set-var`
- `delete-var`
- `power-cycle`
- `power-off`
- `power-on`

In Fuego, boards are defined and configured in a board file, found in the `fuego-ro/boards` directory.

You can use `ftc list-boards` to see a list of the currently configured boards in the Fuego system.

Board attributes (or variables)

Usually, to change the configuration of a board, you manually edit the file for that board and adjust its base settings directly. However, Fuego also allows for viewing board attributes (also referred to as board ‘variables’), and for setting and removing attributes of a board using `ftc` (that is, without having to manually read or edit the board configuration file).

The variables defined in the board configuration file are considered its ‘base’ settings or base attributes. These attributes are considered statically defined for a board. Fuego also allows you to store information about a board that is considered dynamic. This information is stored in a board configuration file in the `fuego-rw/boards` directory.

Also, Fuego automatically assigns certain functions to a board based on the value of the `DISTRIB` variable for the board. These functions are called overlay functions, because they can be overridden (or “overlaid”) with functions from the board configuration file.

`ftc query-board`

You can use the `ftc query-board` command to view any of the configured or calculated information about a board. This includes its base variables, dynamic variables, and overlay functions.

To see all of the attributes of a board, use `ftc query-board` and specify the board to inspect, like this:

```
ftc query-board -b beaglebone
```

The output may be quite verbose. To see just a list of attributes names, (ie without their values), use:

```
ftc query-board -q -b beaglebone
```

To see the value of a single attribute, use the `-n` option, and specify the attribute name:

```
ftc query-board -b beaglebone -n TOOLCHAIN
```

The `set-var` and `delete-var` commands are used to set or delete an individual dynamic variable for a board. These `ftc` commands are intended for programs that automatically configure attributes of a board, and are not usually used by users directly.

`ftc set-var` and `delete-var`

Here are some examples of using `set-var` and `delete-var` on a board:

```
ftc set-var -b beaglebone FOO_COUNT=5  
ftc delete-var -b beaglebone FOO_COUNT
```

These would set `FOO_COUNT` (to the value of ‘5’) in the beaglebone board attributes or remove `FOO_COUNT` from the beaglebone board attributes, respectively.

Finally, `ftc` includes commands for performing power control of a board. When Fuego detects that a board is not responding, it tries to automatically restart the board by doing a power reset.

ftc power commands

The three commands that can be used to manipulate board power are: `power-cycle`, `power-off`, and `power-on`. Here is an example of a power-related command for a board:

```
ftc power-cycle -b beaglebone
```

Note: In order for Fuego to be able to manipulate the power for a board, the board must have a supported `BOARD_CONTROL` system in its configuration.

12.2.4 Jenkins-related commands

These commands are used for interacting with Jenkins, from the command line.

- `add-job(s)`
- `add-node(s)`
- `list-jobs`
- `list-nodes`
- `rm-job(s)`
- `rm-node(s)`
- `build-job(s)`
- `add-view`

By default, Fuego is installed with the Jenkins CI system. Fuego supports integration with many Jenkins operations. This includes `ftc` commands for adding Fuego board and tests to Jenkins, and manipulating those items - listing them, removing them, and in the case of jobs, running them.

Of course, if you are using Fuego in an installation without the Jenkins CI system, none of these commands are relevant, and they may safely be ignored.

As a grammatical courtesy, for some of these commands, you can omit the trailing 's' in the command name, and the command will still work. For example: `ftc add-job` does the exact same thing as `ftc add-jobs`.

When a user wants to install a Fuego test as a job in Jenkins, they use the `ftc add-node` command, to first make sure that the appropriate node (Fuego board) is registered with Jenkins, and then `ftc add-jobs` to add the Fuego tests as jobs within the Jenkins system.

To view or remove nodes or jobs, the `list-nodes`, `list-jobs`, `rm-nodes`, or `rm-jobs` commands are used, respectively.

Finally, the `ftc build-job` command can be used to start a Jenkins job. This is the preferred mechanism to start a Fuego test that has been registered with Jenkins via `ftc add-job`.

ftc add-nodes

`ftc add-nodes` is used to register a Fuego board with the Jenkins interface as an execution node (an object that can run a test).

Once you have added a board to Fuego, you can add it to the Jenkins interface, using:

```
ftc add-node -b beaglebone
```

Usually this will be done once, by the Fuego administrator, when a board is initially added to Fuego. Please see [Adding a Board](#) for instructions to add a new board to Fuego.

ftc add-jobs

The `ftc add-jobs` command is used to configure Jenkins to run Fuego tests, by creating Jenkins job configurations for them. The command provides a few different ways to specify the set of tests to add Jenkins, as well as some options to set other test control options that are used with Fuego when the respective jobs are executed.

The overall usage for `add-jobs` is:

```
ftc add-jobs -b <board>[,board2...] -t <test> [-s <testspec>]
               [--timeout <timeout>] [--rebuild <true|false>] [--reboot <true|false>]
               [--precleanup <true|false>] [--postcleanup <true|false>]
```

And here is an example of a command:

```
Example: ftc add-jobs -b beaglebone -t Dhrystone --timeout 5m --rebuild false
```

This would create the Jenkins job: `beaglebone.default.Benchmark.Dhrystone` (where ‘default’ means the ‘default’ spec (or variant) for this test).

To see a list of possible boards, tests or specs, use `ftc list-boards`, `ftc list-tests` or `ftc list-specs -t <test_name>` respectively.

The other options are used to set the values for the options used with `ftc run-test` when the test is executed by Jenkins.

- `timeout`: integer with a suffix from ‘smhd’ (seconds, minutes, hours, days).
- `rebuild`: if true rebuild the test source even if it was already built.
- `reboot`: if true reboot the board before the test.
- `precleanup`: if true clean the board’s test folder before the test.
- `postcleanup`: if true clean the board’s test folder after the test.

See the section [ftc run-test](#) for more information on the meaning of these options.

Note that you can specify more than one board using a comma-separated list for the `<board>` argument. For example:

```
ftc add-jobs -b board1,board2 -t hello_world
```

If you specify a batch test, then Fuego will scan the list of tests included in that batch test, and create jobs for all of them. For example:

```
ftc add-jobs -b beaglebone -t batch_smoketest
```

will try to create a job for each test referenced in the `batch_smoketest` batch job.

ftc add-view

Finally, Fuego provides a convenience command for easily creating a Jenkins ‘view’. Jenkins supports the ability to organize test jobs by creating views in the user interface. However, it is often convenient to create a view for a small set of Fuego jobs, based on their name.

`ftc add-view` creates a new ‘view’ in Jenkins, with a filter based on the parameter provided.

The syntax for adding a view is:

```
ftc add-view <view-name> [<job_spec>]
```

Basically, you provide a view name, and then an optional string indicating the set of jobs you would like included in that view in the Jenkins dashboard.

You can select individual jobs by name, or use a regular expression (ie with wildcards) to specify the set of jobs to include.

If the job specification starts with “=”, it is a comma-separated list of job names. If not, then it is used as a regular expression.

As a special case, when the command is used without a ‘job_spec’ argument then the view is created with a job_spec consisting of the view-name with wildcards added to the beginning and ending of it.

Here are some examples:

Example 1: `ftc add-view batch ".*.batch"`

Example 2: `ftc add-view network =bbb.default.Functional.ipv6connect,bbb.default.
↪Functional.netperf`

Example 3: `ftc add-view LTP`

Example 3 would add a view with a name of ‘LTP’ and a job specification of “.*LTP.*”. This would result in a view that included all jobs that have “LTP” anywhere in their job names.

12.2.5 Commands for working with a Fuego (or other) server

The Fuego server feature supports executing tests, and sharing test definitions and test run results, between multiple test sites. This feature is currently under construction.

The following commands are related to using Fuego in conjunction with a Fuego server:

- `get-board`
- `get-run`
- `install-run`
- `install-test`
- `list-requests`
- `package-run`
- `package-test`
- `put-binary-package`
- `put-request`
- `put-run`

- put-test
- query-request
- rm-request
- update-board

The following commands support remote operations (using the ‘-r’ or ‘-remote’ flags):

- list-boards
- list-tests
- list-runs

These commands are used for performing operations with a Fuego server. A Fuego server supports registering boards, and storing test packages, binary test packages, run data, and test execution requests. the `package-test` and `package-run` allow for creating packages for a test and a run, respectively. These can be uploaded to the server, or sent directly to another developer, who can install them on their system.

Users can download tests, binary-packaged tests, and runs from a server.

A Fuego server allows a collection of Fuego sites to share tests and test results (runs) with each other. It also allows users to request tests to be executed on boards at another site. These test ‘requests’ can be submitted, viewed, and processed by users interacting with the central server.

Caution: The Fuego server feature is currently under construction. You may experiment with it if you would like, but the features are not robust and the documentation is not finished for it yet. Proceed at your own risk with this feature (and these commands).

12.2.6 Miscellaneous commands

The following commands are for various utility functions, unrelated to the other categories of operations:

- config
- help
- version

The `ftc config` command allow quering the current `ftc` config file, (located in the `fuego-ro\conf` directory). Use `ftc config -l` to get a list of all config items, and `ftc config <config_name>` to get the value of the indicated config item. Usually, humans will not use this, as they can inspect the file manually. The `config` command is intended for use for external tools that want to determine the value for a specific Fuego configuration item.

The `ftc help` command is used to get online usage help for `ftc` and for individual `ftc` commands.

Examples:

```
ftc help - will show a list of all available ftc commands

ftc help list-boards - will show the help for the 'list-boards' command
```

`ftc version` shows the current version of `ftc`.

12.3 ftc run-test

One of the most important commands that ftc can execute is the ‘run-test’ command. This is the command actually used to perform a test on a board. A test can be executed either from the command line (using the ftc command, or it can be executed from Jenkins, via a job definition (which is made with the ftc add-jobs command). Even when running from Jenkins, the ftc run-test command is used to actually execute the test.

When running a test, multiple arguments and options are supported.

Arguments that are required for this are the board name and the test name. The board is specified using ‘-b’ and the test is specified using ‘-t’. foo bar

Here is the ftc run-test usage:

```
Usage: ftc run-test -b <board> -t <test> [-s <spec>] [-p <phases>]
  [--timeout <timeout>]
  [--rebuild <true|false>]
  [--reboot <true|false>]
  [--precleanup <true|false>]
  [--postcleanup <true|false>]
  [--batch]
  [--dynamic-vars <variable assignments or python_dict>]
```

12.3.1 Choosing a test spec

Some Fuego tests include different variants of a test, that can be selected using the test spec. You can see the list of specs for a test using the ftc list-specs command. If no spec is specified for run-test, then the “default” spec is used, which is usually the test executed in its most common configuration.

12.3.2 Test control options

Various other flags control aspects of test execution:

- **timeout**: specify the maximum time the test is allowed to run (default: 30m = 30 minutes)
- **rebuild**: if ‘true’ rebuild the test source even if it was already built. (default: ‘false’)
- **reboot**: if ‘true’ reboot the board before the test. (default: ‘false’)
- **precleanup**: if ‘false’, do not clean up the board’s test folder before the test. (default: ‘true’)
- **postcleanup**: if ‘false’ do not clean up the board’s test folder after the test. (default: ‘true’)
- **batch**: generate a batch id for this test

Each of the boolean test control flags can set be ‘true’ or ‘false’.

The control flags are used to pre-reboot the board being tested, or to prevent or force cleaning up the test directory. Usually, Fuego removes all traces of the test upon test completion. When debugging a test, it is often useful to set --postcleanup to ‘false’, so that Fuego won’t remove the test directory on the board at the end of the test. This allows you to inspect the test materials on the board, or run the test manually. Setting --precleanup to ‘false’ is sometimes useful when you want to avoid the deploy phase. (See [phases](#) below.)

The timeout value is specified as an integer and a suffix (one of s, m, h, or d). The suffix corresponds to one of: seconds, minutes, hours, days. For example, a 10-minute timeout would be specified as --timeout 10m.

The batch id is a number used to group tests together for reporting purposes. if --batch is specified, Fuego will select a new batch id for the test, and set the FUEGO_BATCH_ID environment variable. This will be recorded for this test and

any sub-tests called during execution of the test). You can filter tests using the batch id in a `--where` clause, in the `ftc gen-report` command.

Note: If you would like to specify your own batch id for a test, you can do so by setting the `FUEGO_BATCH_ID` environment variable to your own value before calling `ftc run-test`.

12.3.3 Test variables

It is possible to override one or more test variables on the `ftc run-test` command line, using the `--dynamic-vars` option.

This allows overriding the variables in a test spec from the `ftc` command line. For example, the `Benchmark.Dhrystone` test uses a test variable of `LOOPS` to indicate the number of times to execute the Dhrystone operations. The value of this variable in the default spec for this test is 10000000 (10 million). You could override this value at the command line, using

```
ftc run-test -b bbb -t Dhrystone --dynamic-vars "LOOPS=400000000"
```

You can specify the variable value using a python dictionary expression or using simple `NAME=VALUE` syntax. Here is an example using python dictionary syntax:

```
ftc run-test -b bbb -t Benchmark.Dhrystone --dynamic-vars '{"LOOPS': '4000000000'}"
```

Note that both of these syntaxes allow for multiple variables to be specified. In the case of `NAME=VALUE` pairs, separate the pairs with a comma, like this:

```
--dynamic-vars "VAR1=value1,VAR2=Another value"
```

12.3.4 Phases

A test is normally run in phases, one sequentially after the other. However, in special circumstances (such as when debugging a test during test development) it may be useful to only execute certain phases of the test. Some phases take a long time, and it can be helpful to skip them when debugging a test.

When specifying a set of phases with the `run-test -p` option, each phase is represented by a single character:

```
p = pre_test
c = pre_check
b = build
d = deploy
s = snapshot
r = run
t = post_test
a = processing
m = make binary package
```

To control the phases executed during a test run, use the `-p` option, and specify a list of characters corresponding to the phases you want to execute.

For example:

```
ftc run-test -b myboard -t mytest -p pcbd
```

would run the `pre_test`, `pre_check`, `build` and `deploy` phases of the test ‘mytest’.

This is useful during development of a test (ie. for testing tests). Use caution running later phases of a test without their normal precursors (e.g. specifying to execute the `run` phase, without also specify to execute the `pre_test`, `build` or `deploy` phases). This can lead to undefined behavior.

Warning: It is almost always desirable to run the `pre_test` phase (‘-p’), so use caution omitting that phase from your list. In general, using phase selection is quite tricky, and unless you know what each phase does (and its side effects), it may lead to unexpected results.

12.3.5 Results and Result code

When a test is run, log files and results are placed in a log directory. The log directory is based on a set of attributes for the test run, underneath the `fuego-rw/logs` directory. The pattern for the directory name is:

```
<test_name>/<board>.<spec>.<build_number>.<build_number>
```

So a full log directory name might look like this:

```
fuego-rw/logs/Benchmark.signaltest/beaglebone.default.5.5
```

For more details about the log files that are produced during a test run, see [Log files](#).

In addition to populating the Fuego log directory, if Jenkins is being used and a corresponding Jenkins job is defined for the test, then the Jenkins interface will be populated with information for that test run (referred to as a “build” in Jenkins). A visualization of the test results (for example a chart or a table), may be prepared for display in the Jenkins interface.

Also, although a test may have multiple individual test cases that it executes, the overall status of the test is reported via the return code from `ftc run-test`. This will be 0 for success, and something else for test failure. Usually, a non-zero result will be the value that was returned by the main test program that was run on the board.

LOG FILES

During test execution, Fuego collects several different logs.

These represent different aspects of the test system and the status of the board under test, and are used for different purposes.

Here are the main logs collected or generated during a test:

- **console log**
- **test log**
- **run.json**

In addition to the main test logs and results files, there are other files (test artifacts) with information collected during the test. These are used to debug test execution, board farm status, board status, and results processing.

- **board snapshot**
- **syslogs**
- **devlog**
- **prolog.sh**
- **‘result’ directory**

These are located in the ‘log’ directory (also know as the ‘run’ directory), at: `/fuego-rw/logs/<test_name>/<board>.<spec>.<build_id>.<build_number>`

13.1 Main results logs

These logs are the results of test execution, and have output from different parts of the system.

13.1.1 Console log

The console log is collected by Fuego during the entire execution of the test. This includes output from commands executed during every phase of test execution.

The console log is called: `consolelog.txt`

The main purpose of this log is to show command output during every phase of the test (not just the ‘run’ phase), in order to debug test execution, and show test progress.

The console log is the same thing that shows up in the Jenkins interface during test execution. It is also available after the test finishes, by clicking on the ‘Console Output’ link on an individual Jenkins build page.

Debug data

When debug options are turned on, then each line of the Fuego core and the test's `fuego_test.sh` shell script is output and put into the console log as it is executed.

During test execution, several levels of shell script are run (that is, “source” statements are executed). The number of ‘+’ signs preceding a line in the console log indicates the depth (or invocation nesting level) for that line, in the shell execution.

There are “phase” messages added to the log during the test, which help identify which part of the test a particular sequence is in. This can help you debug what part of a test (`pre_test`, `build`, `deploy`, `run`, `post_test`, etc.) is failing, if any.

You can control which phases emit debugging messages, by setting the `FUEGO_LOG_LEVELS` variable for a test.

This can be done in the Jenkins job config for a job, or by setting the environment variable to an appropriate string before running `ftc run-test`.

13.1.2 Test log

This is the output produced by the test program itself. During execution of a Fuego test, many operations are performed, usually including the execution of a program on the board under test. That program's output is collected and saved. Data for this log is usually only collected during the ‘run’ phase of a test.

Certain Fuego functions, that are used to execute the key commands for a test, save their results to the test log. These include the *report*, *report_append*, *report_live*, and *log_this* functions.

In the case of the *report*, and *report_append* functions, the output from the commands called by these functions is saved into a log file which is retrieved from the target at the end of the test.

The name of the test log for a test is: `testlog.txt`

The purpose of this log is to preserve the actual output of the test program itself.

13.1.3 run.json

`run.json` is a file that has information about the test run, as well as detailed information about test results. The test meta-data and results are expressed in a unified, canonical format, such that results can be queried and compared across test runs (and potentially different test systems). No matter what format output a test program returns, Fuego converts this into a machine-readable list of testcase results and (for benchmarks) test measurement results.

See *run.json* for details of the contents of this file.

13.2 Other test artifacts

Other files are produced during a test and stored in the log directory. Here are descriptions of each of these test artifacts and purpose and expected usage of them.

13.2.1 Board snapshot

The board snapshot file, called ‘machine-snapshot.txt’, contains information about the process environment and board status at the start of the test. By default it includes the following information:

- shell environment variables for processes executed by the test
- kernel version information (called the ‘firmware revision’)
- process load
- memory status
- mounted filesystems
- running process list
- the status of interrupts

The main purpose of this is to allow the user to see if there were any unusual conditions on the machine, at the time of the test. This is intended to allow diagnosing test failures caused by machine state.

13.2.2 Syslogs

The syslogs record messages from the board’s system log. There are two of these recorded, one before the start of test execution (called the “before” log), and one after test execution (called the “after” log).

These logs include messages from the kernel (if a logger is transferring the messages from the kernel to the system log), as well as messages from programs running on the system, that output to their status to the system log.

The names of the syslogs for a test are:

- `syslog.before.txt`
- `syslog.after.txt`

The purposes of these logs is to capture the condition of the machine before the test started, and also to allow determining any error conditions that were logged on the system while the test was running.

In particular, the difference between the ‘after’ syslog and the ‘before’ syslog is examined by Fuego, to check for any kernel oopses (kernel failures reported to the log). If the kernel has an “oops” during a test, then the test is reported as having failed, regardless of other conditions (test program return code or criteria processing of testcase results).

13.2.3 Devlog

This is a summarized list of board management and control operations performed during the execution of a test.

The name of the devlog for a test run is `devlog.txt`

The main purpose of this log is for debugging the internal Fuego command sequences, and for determining the health of the board within a board farm. Most Fuego end users can ignore this log.

13.2.4 prolog.sh

The prolog.sh is a shell script that is generated during test execution. It contains the Fuego-generated test variables for the test, as well as the overlay functions that were used for this test run.

It is used during test execution, but is left in the log directory to allow for debuggin test execution. Most Fuego users should not need to examine this file. It is useful mainly for developers of the Fuego test system itself, and developers of Fuego tests.

13.2.5 ‘Result’ directory

Some tests produce a directory in the log area called ‘result’. This is automatically produced by some parsers, when they process the testlog.txt file. Sections of test testlog.txt that contain diagnostics information for individual testcases are stored as individual files, so that they can be displayed as separate pages in the Jenkins user interface. For some tests, these individual files are also processed into test-specific reports (specifically, LTP has additional spreadsheet-creation facilities that use the individual files in the ‘results’ directory.

These directories can safely be ignored by the user in most cases.

13.3 Summary

In the Fuego version 1.5 (released in 2019), the log directories are as follows:

- Fuego logs:
 - /fuego-rw/logs/<testname>/<board>.<spec>.<build_id>.<build_number>
 - * consolelog.txt - collection of all output from all commands during a test
 - * testlog.txt - the output of the actual test program
 - * run.json - test meta-data and results in canonical format
 - * machine-snapshot.txt - status of board before test
 - * syslog.before.txt - system log of board under test, saved before test execution
 - * syslog.after.txt - system log of board under test, saved after test execution
 - * devlog.txt - list of board command and control operations
- jenkins files:
 - /var/lib/jenkins/jobs/<jobname>/builds/buildnum/
 - * build.xml - Jenkins meta-data about the build
 - * log - same as the console log above
- per-test data files:
 - /fuego-rw/logs/<testname>
 - * flat_plot_data.txt - has results data in “flat” ASCII text format
 - * flot_chart_data.json - has chart data in json and HTML format

GENERATING REPORTS

Usually, test results from Fuego tests are examined in the Jenkins User interface. However, Fuego also provides the capability to generate reports outside of Jenkins.

A report consists of data from one or more test runs, collected and presented in the report in an organized fashion. The Fuego system includes the capability to generate lists of test results, providing a number of options to control the content, format, and location of reports.

The command used to generate reports in Fuego is `ftc gen-report`

14.1 ftc gen-report

The `ftc gen-report` command gives you control over the results that are reported, as well as the content (exact fields and headers) and format of the report.

Using `ftc gen-report`, you can:

- select the test runs from which to report results
- select the header fields to show in the report
- select the data fields to show in the report
- filter the data by results (for example to show only failures)
- select the format of the report
- select the location for the report output file

14.1.1 Overall command usage

To generate a report, you use `ftc gen-report` and use command line options to control the data and format of the report.

The overall command usage is:

```
ftc gen-report [filter-options] [field-options] [format-option]
```

The online command usage (obtainable with `ftc gen-report help` is shown in the section [ftc gen-report usage help](#) below).

Details about the options for generating reports is shown in the following sections.

14.1.2 Selecting data elements to include in the report

A report consists of data from one or more test runs, collected and presented in the report.

Each report has a header section and a body section.

Header section

The fields in the header section are meta-data about the tests included in the report, such as the set of test names, boards that the tests were run on, the kernel versions tested, the `start_time` of the tests, and the date that the report was generated.

You can control which fields are shown in the header section of the report, using the `--header-fields` option.

The list of header fields available is:

- `test`
- `board`
- `kernel`
- `timestamp`
- `report_date`

Header fields that have multiple values will be shown as a comma-separated list (like the board names), or as a range of values with a start and end (e.g. for the timestamp field).

If no `--header-fields` option is used, then the default set of header fields displayed in the header section of the report (and their field order), is:

`test,board,kernel,timestamp,report_date`

Data section

The data section of the report has data about test results for the selected runs. This includes information such as the test names, board, testcase names, and results.

The data in the report is obtained from a set of test runs that are selected for the report (see below). The data consists of information about the test run (including, importantly, the status of each test), and can also contain detailed information about individual testcase results.

You can control which fields are shown (and the order they are shown in) in the data section of the report, using the `--fields` option.

The list of fields available for inclusion in the report is:

- `test_name` - full test name
- `test` - short test name
- `spec` - test spec used for the run
- `board` - board that test ran on
- `kernel` - kernel that test ran on
- `timestamp` - time and date of test start
- `start_time` - test start time (in seconds since the epoch)
- `tguid` - test case identifier

- `tguid:result` - test case result
- `status` - test result (the summary result for the whole test, after the test criteria is applied)
- `duration_ms` - the duration of the test in milliseconds

The order of the fields in each report line corresponds to the order in which the fields are specified as arguments to the `--fields` option. If no `--fields` option is used, then the default set of fields (and their order), is:

```
test_name,spec,board,timestamp,tguid,tguid:result
```

Difference between ‘status’ and ‘tguid:result’

For each test displayed in the report, the test’s status may be shown. The ‘status’ field for a test run is the overall result of the entire test. Only one of these is shown per test run.

Depending on the fields displayed in the report, a report may also display individual testcase names and results. In the report generator, a testcase name is called a ‘tguid’. This stands for “testcase globally unique identifier”. Each tguid consists of a string containing one or more elements indicating the test suite name, test set name, test case name and measurement name (in the case of benchmark measurement results)

The result for each test case is one of PASS, FAIL, SKIP or ERROR. The result for a measurement is always a number.

An example of a tguid for a measurement is:

```
IOzone.2048_Kb_Record_Write.Random_write.Score
```

The associated tguid for this particular test case is:

```
IOzone.2048_Kb_Record_Write.Random_write
```

These indicate that this test case and measurement are for the IOzone test suite, the 2048_KB_Record_Write test set, and the Random_write test case within that set. The actual benchmark measurement for this test case is named its ‘Score’ by this test.

A single Fuego test will only have a single ‘status’, but may have multiple test cases (tguid) that can appear in a report.

14.1.3 Selecting data to include in the report

You can filter the set of runs (and test cases within those runs) to include in the report by using the `--where` option. The string following the `--where` option is called a “where clause”. It is used to specify a condition that must be met for a run (or a piece of test result data) to be included in the report. This is similar in functionality to the WHERE used in the SQL database query language, but with different syntax.

Multiple where clauses can be specified for a single report, either as a comma-separated list following a single `--where` option, or via separate `--where` options. (See below for some examples.)

Where clauses are used to select the data that will be included in the report. This includes filtering the set of runs whose data will be included, as well as individual testcase items or results that will be included. For example, you can use a where clause to select test runs by board name, test name, or the ‘spec’ that was used for a run. You can also use where clauses to select tests by start time, batch_id, or build_number. You can filter by the result for a full test (known as it’s “status”), as well as filter the data from each test by individual testcase names, and testcase results (both strings and numbers).

In order for an element to be included in the report (a run or a result), it must match all specified where clauses. That is, the filter is an AND operation between all the where clauses that are specified.

where clause syntax

A single ‘where clause’ consists of a field_name, an operator, and a value, like this:

```
board=beaglebone
```

Allowed field names are:

- **test** - the full or shortened test name
- **type** - the test type (value may be either “Functional” or “Benchmark”)
- **spec** - the variant or ‘spec’ used for the run
- **board** - the board
- **start_time** - the time when the test was started
- **batch_id** - a batch id string for the run
- **status** - the overall result of a test run
- **build_number** - the build number for the run
- **tguid** - an individual testcase identifier (globally unique id)
- **tguid:result** - an individual testcase result

Allowed operators are:

- ‘=’ - the field has a value equal to the specified value
- ‘<’ - the field has a value less than the specified value
- ‘<=’ - the field has a value less than or equal to the specified value
- ‘>’ - the field has a value greater than the specified value
- ‘>=’ - the field has a value greather than or equal to the specified value
- ‘!=’ - the field has a valud that is not equal to the specified value
- ‘=~’ - the field has a value that matches the specified regular expression

The ‘=~’ (match) operator is used with a regular expression that the field value must match, in order for the run (or data element) to be included in the report. The regular expression is expressed in python ‘re’ syntax. See <https://docs.python.org/2/library/re.html#regular-expression-syntax> for information about the regular expression syntax supported with the ‘=~’ (match) operator.

Here are some example where clauses:

```
example 1: --where test=LTP
example 2: --where board=beaglebone,test=bonnie
example 3: --where board=beaglebone --where test=bonnie
example 4: --where "start_time>2 hours ago"
example 5: --where batch_id=12
example 6: --where "tguid=~.*udp.*"
```

(continues on next page)

(continued from previous page)

```
example 7: --where tguid:result=FAIL
```

```
example 8: --where test=fio,"tguid:result>10000"
```

Here are descriptions of these examples:

- example 1 says to generate a report of LTP test runs
- example 2 says to show test results for the ‘bonnie’ test on the ‘beaglebone’ board
- example 3 says the same thing as example 2. Examples 2 and 3 are different ways of expressing multiple where clauses for a single report generation command. Their effect is identical.
- example 4 says to show tests started in the last 2 hours
- example 5 says to show tests with batch_id=12
- example 6 says to show tests results where the testcase identifier includes the string ‘udp’
- example 7 says to show test results that failed.
- example 8 says to show test results for the ‘fio’ test, where the benchmark result was greater than 10,000.

Here are some additional tips for specifying where clauses:

- When using a where clause that has spaces, enclose the value, or the whole expression in quotes, to avoid the shell splitting the command line argument
- When using ‘>’ or ‘<’, enclose the expression in quotes, to avoid the shell interpreting the comparison operators as output or input redirection, respectively
- When using shell wildcard characters (such as ‘*’ or ‘?’) in a ‘=~’ regular expression, enclose the expression in quotes, to avoid the shell interpreting the wildcards as file-matching strings.

start_time tricks

It can be a bit tricky to specify the right start_time for the where clause. Say, for example, you want to get the report for tests within the last few hours. This is possible; the parsing for start_time in the where clause is flexible. But it may be confusing.

In general, Fuego supports specifying the start_time value in a where clause using English strings that specify absolute or relative time. For example, the following strings are supported:

- --where “start_time>2023-02-24 10:15”
- --where “start_time>today at 10:00”
- --where “start_time>oct 1 2022”
- --where “start_time>1 hour ago”

The last of these examples uses a phrase that has relative time sense. That is, the time is expressed relative to the time of invocation of the `ftc`.

In Fuego, start_time comparison operations are in absolute time sense, not in relative time sense, even if the time value is expressed using relative wording. Thus

```
--where "start_time>1 hour ago"
```

will include tests started “within the last hour”. That is: show tests with a start_time time that has a value greater than the time 1 hour before the `ftc gen-report` command was run).

The where clause:

```
--where "start_time>2023-02-24 00:01"
```

means include tests with a start time after 12:01 am Feb 24, 2023.

As a help, if you use the `-v` option with `ftc gen-report`, Fuego will tell you in human-readable form what the `start_time` value is that Fuego interpreted for your where clause.

Here's an example:

```
$ ftc -v gen-report --where "start_time>last year"
start_time value in 'where' is '2022-01-01 09:00:00.000002'
=====
      **** Fuego Test Report ****
...

```

Using `batch_id` to group runs for a report

Another way to group test runs for a report is by using a `'batch_id'`.

When you run a batch test in Fuego, all the individual tests in the batch are assigned the same `batch_id`. This batch id is printed during the test run, or you can use `ftc gen-report` with the `--fields` option to see the batch-id for any test or set of tests.

Using a `'batch-id'` makes it easy to generate a report for a collection of tests, by using a where clause that specifies the `batch_id`.

Even if you are not using a Fuego batch test (where a `batch_id` is assigned automatically), you can manually assign a batch-id to a collection of tests, by setting the variable `FUEGO_BATCH_ID` to some unique string, before executing the `ftc run-test` command line, for a set of tests.

Here is an example of using a custom `'batch_id'` for a collection of tests, to enable generating a report for that collection of test runs:

```
export FUEGO_BATCH_ID=CI-loop-26
ftc run-test -b beaglebone,minnowboard -t fuego_board_check
...
ftc run-test -b beaglebone,raspberry-pi -t fio
...
ftc run-test -b minnowboard -t cyclicttest
...

ftc gen-report --where batch_id=CI-loop-26
...

```

Filtering the data by results

You can also filter the data for a test report by specifying where clauses that specify matches or comparisons for test results (status), individual testcase results or Benchmark test measurement results.

A common report filter is one where only errors, skips and failures are reported, and PASS results are ignored (that is, omitted from the report). This can be done by filtering on the field 'tguid:result'

Here is an example:

```
ftc gen-report --where "start_time>last week" --where tguid:result!=PASS
```

This would generate a report with a list of all failing tests and testcase results, in the last week.

If you are not interested in ERROR or SKIP results, you might use a more specific result filter, looking only at FAIL results.

```
ftc gen-report --where "start_time>last week" --where tguid:result=FAIL
```

Or, you might only want a summary of tests that failed in the last week, ignoring individual testcase results within those tests. To do that, you need to specify a custom field list that does not include tguid:result:

```
ftc gen-report --where "start_time>last week" --where status=FAIL \
  --fields test,spec,board,build_number,timestamp,status
```

14.1.4 Specifying the format and location for the report

Fuego supports outputting the report data in a few different formats.

By default, the report is output in text format, to the stdout of the 'ftc' command. To specify a different format for the report, use the --format option.

The following output formats are supported:

- **txt** - plain text
- **html** - HTML format
- **rst** - reStructured text
- **pdf** - PDF file format
- **excel** - Excel spreadsheet format
- **csv** - comma separated values

The default format, if none is specified with the --format option, is 'txt' (plain text).

Each of the other formats are well-known document formats.

Extra text report format options

A few extra options are available to control the formatting of the output, when using the ‘txt’ output format.

- `-f` - use fixed column widths in the text output format
- `-q` - quiet mode: omit the header section and headings in text output format

By default, when outputting data in the text output format, Fuego will automatically adjust the column widths according to the sizes of the strings in that column. However, it may be desirable to present the values at fixed offsets in each line, so that the data may be more easily parsed by other tools. To use a fixed width of 20 characters for each column (starting at 2nd character of each line), use the ‘`-f`’ option. Here is an example of report output, without the ‘`-f`’:

```
$ ftc gen-report --where board=bbb,test=bc,build_number=1
=====
      **** Fuego Test Report ****
test           : bc
board          : bbb
kernel         : 4.4.155-ti-r155
timestamp      : 2022-03-09_19:07:49
report_date    : 2023-03-15_03:12:19
=====
-----
test_name      spec      board timestamp      tguid tguid:result
-----
Functional.bc  default  bbb   2022-03-09_19:07:49  bc    PASS
-----
```

And here is an example of same report, *with* the ‘`-f`’ option:

```
$ ftc gen-report -f --where board=beaglebone,test=bc,build_number=1
=====
      **** Fuego Test Report ****
test           : bc
board          : bbb
kernel         : 4.4.155-ti-r155
timestamp      : 2022-03-09_19:07:49
report_date    : 2023-03-15_03:12:04
=====
-----
↪ test_name      spec      board      timestamp      ↪
↪ tguid          tguid:result
-----
↪ Functional.bc  default  bbb      2022-03-09_19:07:49  bc ↪
↪                PASS
-----
↪ -----
```

You can get more concise output, in text mode, using the ‘`-q`’ option.

When you use ‘`-q`’ with `ftc gen-report` and you are in text mode, only the data section of the report is shown. The header section is omitted, and the heading for the rows of data are also omitted. Only the fields for each line of data in the report are printed. Also, the report data starts in column 0, instead of column 2. (That is, the report data is not indented by 2 spaces).

The purpose of ‘quiet mode’ (the ‘-q’ option) in report generation is to allow you to extract data from the Fuego test results, in a format that can easily be parsed by other tools.

Here is an example:

```
$ ftc gen-report -q --where board=beaglebone,build_number=1,tguid=Read.Seq.speed \  
  --fields tguid:result  
21457.00
```

Output location

Fuego can output report data to stdout (the output of the `ftc` command), or save the data to a file. The filename for a report sent to a file follows the pattern: “Test_report_<date_and_time>.<ext>”.

The default output location depends on the format used for the report. The output directory can be changed using the `-o` option.

Here is the default output location for the different output format types:

- `txt` - output to stdout
- `html` - output to stdout
- `rst` - output to stdout
- `csv` - output to `/fuego-rw/reports/Test_report_<date-and-time>.csv`
- `excel` - output to `/fuego-rw/reports/Test_report_<date-and-time>.xls`
- `pdf` - output to `/fuego-rw/reports/Test_report_<date_and_time>.pdf`

The ‘`txt`’, ‘`html`’ and ‘`rst`’ formats are output to stdout, unless a report directory is specified (using the `-o` option). The ‘`csv`’, ‘`excel`’, and ‘`pdf`’ formats are written to a file in the Fuego reports directory. By default the report directory is `/fuego-rw/reports`, inside the docker container. This directory is visible on the host in the `fuego-rw/reports` directory where you installed Fuego.

You can use “`-o -`” to use the default report directory for ‘`txt`’, ‘`html`’ and ‘`rst`’ formats. Instead of writing the report to stdout, the Fuego will write reports using these formats to `fuego-rw/reports`, using the standard report filename and an appropriate filename extension.

Important: When using a Fuego container, the report directory that is used with the `ftc gen-report` command will be relative to the root of the container, not the host filesystem. That is, if you run: `ftc gen-report -o /tmp`, from outside the container, the report file will be placed in the `/tmp` directory inside the container, *NOT* the `/tmp` directory on the host.

This may cause confusion. When the Fuego reports directory (`fuego-rw/reports`) is used, this distinction is not a problem since the `fuego-rw` directory is visible in both the container and the host.

14.2 ftc gen-report usage help

Here is output generated when you type `ftc gen-report help`:

`ftc gen-report`: Generate a report **from** a **set** of runs

```
Usage: ftc gen-report [--where <where-clause1>[,<where-clausen>]...] \
    [-f] [-q]
    [--format [txt|html|pdf|excel|csv|rst]] \
    [--header_fields <field_list>] \
    [--fields <field_list>] \
    [--layout <report_name>]
    [-o <report_dir>]
```

Generates a report **from** **test** run data **as** specified. The `where` option controls which runs are included **in** the report. The `format` option controls the text encoding of the report, **and** the layout specifies a report style. If no arguments are provided, the defaults of "all tests", "txt" output, **and** a layout consisting of summary results **is** used.

`-f` use fixed column widths **in** the text output **format**
`-q` omit table header **and** column headings **in** text output **format**

txt, html **and** rst formats are output to stdout, unless a report **dir is** specified. pdf, excel **and** csv formats are written to a file **in** the report directory. By default the report directory **is** `/fuego-rw/reports`, but this can be overridden **with** the `-o` option. Use "`-o -`" to use the default report directory **for** txt, html **and** rst formats.

The `--where` option can be used to specify one **or** more 'where' specifiers to **filter** the **list** of runs. Each where clause **is** separated by a comma. A 'where clause' consists of a field_name, an operator **and** a value. Allowed field names are: test, type, spec, board, start_time, result, batch_id, status, build_number, tguid, **and** tguid:result. Allowed operators are: '=', '<', '<=', '>', '>=', '!=', '~='. The '~= ' operator means the value **is** a regular expression to match, **for** the indicated field. Here are some example where options:

```
--where test=LTP
--where test=bonnie,board=beaglebone
--where "start_time>2 hours ago"
--where batch_id=12
--where tguid=~.*udp.*
--where tguid:result=FAIL
```

The `--header_fields` **and** `--fields` options allow specifying a **list** of field names (comma-separated) **for** inclusion **in** the header **and** body of the report, respectively. The default field lists, **if** none are specified are: (**for** headers) test,board,kernel,timestamp,report_date **and** (**for** fields) test_name,spec,board,timestamp,tguid,tguid:result.

Here **is** an example:

```
ftc gen-report --header_fields test --fields timestamp,tguid,tguid:result
```

ADDING OR CUSTOMIZING A DISTRIBUTION

15.1 Introduction

Although Fuego is configured to execute on a standard Linux distribution, Fuego supports customizing certain aspects of its interaction with the system under test. Fuego uses several features of the operating system on the board to perform aspects of its test execution. This includes things like accessing the system log, flushing file system caches, and rebooting the board. The ability to customize Fuego's interaction with the system under test is useful in case you have a non-standard Linux distribution (where, say, certain features of Linux are missing or changed), or when you are trying to use Fuego with a non-Linux system.

A developer can customize the distribution layer of Fuego in one of two ways:

- adding overlay functions to a board file by creating a new
- distribution overlay file

15.2 Distribution overlay file

A distribution overlay file can be added to Fuego, by adding a new “.dist” file to the directory: `fuego-core/overlays/distrib`

The *distribution* functions are defined in the file: `fuego-core/overlays/base/base-distrib.fuegoclass` These include functions for doing certain operations on your board, including:

- `ov_get_firmware`
- `ov_rootfs_reboot`
- `ov_rootfs_state`
- `ov_logger`
- `ov_rootfs_sync`
- `ov_rootfs_drop_caches`
- `ov_rootfs_oom`
- `ov_rootfs_kill`
- `ov_rootfs_logread`

You can define your own *distribution* overlay by defining a new “.dist” file in `fuego-core/overlays/distributions`. (e.g. `my-dist.dist`) Basically, you inherit functions from `base-distrib.fuegoclass`, and write override functions in `mydist.dist` to perform those operations the way they need to be done on your distribution.

You can look up what each override function should do by reading the fuegoclass code, or looking at the function documentation at: [Test Script APIs](#)

The inheritance mechanism and syntax for Fuego overlay files is described at: [Overlay Generation](#)

The goal of the distribution abstraction layer in Fuego is to allow you to customize Fuego operations to match what is available on your target board. For example, the default (base class) `ov_rootfs_logread()` function assumes that the target board has the command `"/sbin/logread"` that can be used to read the system log. If your distribution does not have `"/sbin/logread"`, or indeed if there is no system log, then you would need to override `ov_rootfs_logread()` to do something appropriate for your distribution or OS.

Note: In fact, this is a common enough situation that there is already a 'nologread.dist' file already in the overlay/distrib directory.

Similarly, `ov_rootfs_kill` uses the `/proc` filesystem, `/proc/$pid/status`, and the `cat`, `grep`, `kill` and `sleep` commands on the target board to do its work. If our distribution is missing any of these, then you would need to override `ov_rootfs_kill()` with a function that did the appropriate thing on your distribution (or OS).

15.2.1 Existing distribution overlay files

Fuego provides a few distribution overlay files for certain situations that commonly occur in embedded Linux testing.

- `nologread.dist` - for systems that do not have a 'logread' command
- `nosyslogd.dist` - for systems that don't have any system logger

15.3 Referencing the distribution in the board file

Inside the board file for your board, indicate the distribution overlay you are using by setting the `DISTRIB` variable.

If the `DISTRIB` variable is not set, then the default distribution overlay functions are used.

For example, if your embedded distribution of Linux does not have a system logger, you can override the normal logging interaction of Fuego by using the 'nosyslogd.dist' distribution overlay. To do this, add the following line to the board file for target board where this is the case:

```
DISTRIB="nosyslogd.dist"
```

15.4 Testing Fuego/distribution interactions

There is a test you can run to see if the minimal command set required by Fuego is supported on your board. It does not require a toolchain, since it only runs shell script code on the board. The test is `Functional.fuego_board_check`.

This test may work on your board, if your board supports a POSIX shell interface. However, note that this test reflects the commands that are used by Fuego core and by the default distribution overlay. If you make your own distribution overlay, you may want to create a version of this test that omits checks for things that your distribution does not support, or that adds checks for things that your distribution overlay uses to interact with the board.

15.5 Notes

Fuego does not yet fully support testing non-Linux operating systems. There is work-in-progress to support testing of NuttX, but that feature is not complete as of this writing. In any event, Fuego does include a ‘NuttX’ distribution overlay, which may provide some ideas if you wish to write your own overlay for a non-Linux OS.

15.5.1 NuttX distribution overlay

By way of illustration, here are the contents of the NuttX distribution overlay file (fuego-core/overlays/distribs/nuttX.dist).

```

override-func ov_get_firmware() {
    FW="$(cmd uname -a)"
}

override-func ov_rootfs_reboot() {
    cmd "reboot"
}

override-func ov_init_dir() {
    # no-op
    true
}

override-func ov_remove_and_init_dir() {
    # no-op
    true
}

override-func ov_rootfs_state() {
    cmd "echo; date; echo; free; echo; ps; echo; mount" || \
        abort_job "Error executing rootfs_status operation on target"
}

override-func ov_logger() {
    # messages are in $@, just emit them
    echo "Fuego log messages: $@"
}

# $1 = tmp dir, $2 = before or after
override-func ov_rootfs_logread() {
    # no-op
    true
}

override-func ov_rootfs_sync() {
    # no-op
    true
}

override-func ov_rootfs_drop_caches() {
    # no-op

```

(continues on next page)

(continued from previous page)

```
        true
    }

    override-func ov_rootfs_oom() {
        # no-op
        true
    }

    override-func ov_rootfs_kill() {
        # no-op
        true
    }
}
```

15.5.2 Hypothetical QNX distribution

Say you wanted to add support for testing QNX with Fuego.

Here are some first steps to add a QNX distribution overlay:

- set up your board file
- create a custom QNX.dist (stubbing out or replacing base class functions as needed)
 - you could copy null.dist to QNX.dist, and deciding which items to replace with QNX-specific functionality
- add `DISTRIB="QNX.dist"` to your board file
- run the `Functional.fuego_board_check` test (using `ftc`, or adding the node and job to Jenkins and building the job using the Jenkins interface), and
- examine the console log to see what issues surface

BUILDING DOCUMENTATION

As of July, 2020, the Fuego documentation is currently available in 3 places:

- The `fuego-docs.pdf` generated from TEX files in the `fuego/docs/source` directory
- The FuegoTest wiki, located at: <https://fuegotest.org/wiki/Documentation>
- A set of html files in `fuego/docs/_build/html` that are generated from `.rst` files in `fuego/docs/rst_src`

The `fuego-docs.pdf` file is a legacy file that is several years old. It is only kept around for backwards compatibility. It might be worthwhile to scan it and see if any information is in it that is not in the wiki and migrate it to the wiki.

The fuegotest wiki has the currently-maintained documentation for the project. But there are several issues with this documentation:

- the wiki used is proprietary and slow
- the information is not well-organized
- the information is only available online, as separate pages
 - there is no option to build a single PDF, or use the docs offline
- there is a mixture of information in the wiki
 - not just documentation, but a crude issues tracker, random technical notes testing information, release information and other data that should not be part of official documentation

The `.rst` files are intended to be the future documentation source for the project.

16.1 building the outdated PDF

To build the outdated PDF, cd to `fuego/docs`, and type

```
$ make fuego-docs.pdf
```

This will use latex to build the file `fuego/docs/fuego-docs.pdf`

16.2 building the RST docs

The RST docs can be build in several different formats, including text, html, and pdf. You can type ‘make help’ to get a list of the possible build targets for this documentation. Output is always directed to a directory under fuego/docs/_build.

Here are some of the most popular targets:

16.2.1 html

```
$ make html
```

Documentation will be in fuego/docs/_build/html

The root of the documentation will be in index.html

16.2.2 singlehtml

```
$ make singlehtml
```

Documentation will be in fuego/docs/_build/singlehtml

The complete documentation will be in a single file: index.html (with images and other static content in fuego/docs/_build/singlehtml/_static

16.2.3 latexpdf

```
$ make latexpdf
```

Documentation will be in fuego/docs/_build/latexpdf/Fuego.pdf

FUEGO DEVELOPER NOTES

This page has some detailed notes about Fuego, Jenkins and how they interact.

17.1 Resources

Here are some pages in this wiki with developer information:

- Coding style
- Core_interfaces
- Glossary
- Fuego test results determination
- Fuego_naming_rules
- Fuego Object Details
- Integration with ttc
- Jenkins User Interface
- Jenkins Plugins
- License And Contribution Policy
- Log files
- Metrics
- Overlay Generation
- ovgen feature notes
- Parser module API
- Test Script APIs
- Test package system
- Test server system
- Transport notes
- *Variables*

17.2 Notes

17.2.1 specific questions to answer

What happens when you click on the “run test” button:

- what processes start on the host
 - `java -jar /home/jenkins/slave.jar`, executing a shell running the contents of the job.xml “hudson.tasks.Shell/command” block:
 - * this block is labeled: “Execute shell: Command” in the “Build” section of the job, in the configure page for the job in the Jenkins user interface.
- what interface is used between the executing test (ultimately a bash shell script) and the jenkins processes
 - stop is performed by using the Jenkins REST API – by accessing “<http://localhost:8090/...stop>”
 - see Fuego-Jenkins

Each Jenkins node is defined in Jenkins in: `/var/lib/jenkins/nodes/config.xml`

- The name of the node is used as the “Device” and “NODE_NAME” for a test.
 - These environment variables are passed to the test agent, which is always “`java -jar /home/jenkins/slave.jar`”
- Who calls `ovgen.py`? The core does, at the very start of `main.sh` via the call to `set_overlay_vars` (which is in `overlays.sh:w`)

Jenkins calls:

- `java -jar /fuego-core/slave.jar`
 - with variables:
 - * Device
 - * Reboot
 - * Rebuild
 - * Target_PreCleanup
 - * Target_PostCleanup
 - * TESTDIR
 - * TESTNAME
 - * TESTSPEC
 - * FUEGO_DEBUG
- the “Test Run” section of the Jenkins job for a test has a shell script fragment with the following shell commands:

```
#logging areas=pre_test,pre_check,build,makepkg,deploy,snapshot,run,
#    post_test,processing, parser,criteria,charting
#logging levels=debug,verbose,info,warning,error
#export FUEGO_LOGLEVELS="run:debug,parser:verbose"
export FUEGO_CALLER="jenkins"
ftc run-test -b $NODE_NAME -t Functional.hello_world -s default \
    --timeout 6m \
    --reboot false \
```

(continues on next page)

(continued from previous page)

```
--rebuild false \
--precleanup true \
--postcleanup true
```

Some Jenkins notes:

Jenkins stores its configuration in plain files under JENKINS_HOME. You can edit the data in these files using the web interface, or from the command line using manual editing (and have the changes take affect at runtime by selecting “Reload configuration from disk”).

By default, Jenkins assumes you are doing a continuous integration action of “build the product, then test the product”. It has default support for Java projects.

Fuego seems to use distributed builds (configured in a master/slave fashion).

Jenkins home has (from 2007 docs):

- config.xml - has stuff for the main user interface
- *.xml
- fingerprints - directory for artifact fingerprints
- jobs
 - <JOBNAME>
 - * config.xml
 - * workspace
 - * latest
 - * builds
 - <ID>
 - build.xml
 - log
 - changelog.xml

The docker container interfaces to the outside host filesystem via the following links:

- /fuego-ro -> <host-fuego-location>/fuego-ro
- /fuego-rw -> <host-fuego-location>/fuego-rw
- /fuego-core -> <host-fuego-core-location>

What are all the fields in the “configure node” dialog: Specifically:

- where is “Description” used? - don’t know
- what is “# of executors”? - don’t know for sure
- how is “Remote root directory” used?
 - this is a path inside the Fuego container. I’m not sure what Jenkins uses it for.
- what are Labels used for?
 - as tags for grouping builds
- Launch method: Fuego uses the Jenkins option “Launch slave via execution of command on the Master” The command is “java -jar /fuego-core/slave.jar”

- NOTE: slave.jar comes from fuego-core git repository, under slave.jar

The fuego-core directory structure is:

```
overlays - has the base classes for fuego functions
  base - has core shell functions
  testplans - has json files for parameter specifications (deprecated)
  distribs - has shell functions related to the distro
scripts - has fuego scripts and programs
(things like overlays.sh, loggen.py, parser/common.py, ovgen.py, etc.
slave.jar - java program that Jenkins calls to execute a test
tests - has a directory for each test
  Benchmark.foo
    Benchmark.foo.spec
    foo.sh
    test.yaml
    reference.log
    parser.py
  Functional.bar
  LTP
  etc.
```

What is groovy:

- an interpreted language for Java, used by the scriptler plugin to extend Jenkins

What plugins are installed with Jenkins in the JTA configuration?

- Jenkins Mailer, LDPA, External Monitor Job Type, PAM, Ant, Javadoc
- Jenkins Environment File (special)
- Credentials, SSH Credentials, Jenkins SSH Slags, SSH Agent
- Git Client, Subversion, Token Macro, Maven Integration, CVS
- Parameterized Trigger (special)
- Git, Groovy Label Assignment Extended Choie Parameter
- Rebuilder...
- Groovy Postbuild, ez-templates, HTML Publisher (special)
- JTA Benchmark show plot plugin (special)
- Log Parser Plugin (special)
- Dashboard view (special)
- Compact Columns (special)
- Jenkins Dynamic Parameter (special)
- flot (special) - benchmark graphs plotting plug-in for Fuego

Which of these did Cogent write?

- the flot plugin (not flot itself)

What scriptler scripts are included in JTA?

- getTargets
- getTestplans

- getTests

What language are scriptler scripts in?

- Groovy

What is the Maven plugin for Jenkins?

- Maven is an apache project to build and manage Java projects
 - I don't think the plugin is needed for Fuego

Jenkins refers to a "slave" - what does this mean?

- it refers to a sub-process that can be delegated work. Roughly speaking, Fuego uses the term 'target' instead of 'slave', and modifies the Jenkins interface to support this.

17.2.2 How the tests work

A simple test that requires no building is Functional.bc

- the test script and test program source are found in the directory: `/home/jenkins/tests/Functional.bc`

This runs a shell script on target to test the 'bc' program.

Functional.bc has the files:

```
bc-script.sh
declares "tarball=bc-script.tar.gz"
defines shell functions:
  test_build - calls 'echo' (does nothing)
  test_deploy - calls 'put bc-device.sh'
  test_run - calls 'assert_define', 'report'
             report references bc-device.sh
  test_processing - calls 'log_compare'
                  looking for "OK"
sources $JTA_SCRIPTS_PATH/functional.sh
bc-script.tar.gz
bc-script/bc-device.sh
```

Variables used (in bc-script.sh):

```
BOARD_TESTDIR
TESTDIR
FUNCTIONAL_BC_EXPR
FUNCTIONAL_BC_RESULT
```

A simple test that requires simple building:

Functional.synctest

This test tries to call fsync to write data to a file, but is interrupted with a kill command during the fsync(). If the child dies before the fsync() completes, it is considered success.

It requires shared memory (shmget, shmat) and semaphore IPC (semget and semctl) support in the kernel.

Functional.synctest has the files:

```
synctest.sh
declares "tarball=synctest.tar.gz"
```

(continues on next page)

(continued from previous page)

```
defines shell functions:
    test_build - calls 'make'
    test_deploy - calls 'put'
    test_run - calls 'assert_define', hd_test_mount_prepare, and 'report'
    test_processing - calls 'log_compare'
    looking for "PASS : sync interrupted"
sources $JTA_SCRIPTS_PATH/functional.sh
sync_test.tar.gz
    sync_test/sync_test.c
    sync_test/Makefile
sync_test_p.log
    has "PASS : sync interrupted"
```

Variables used (by sync_test.sh)

```
CFLAGS
LDFLAGS
CC
LD
BOARD_TESTDIR
TESTDIR
FUNCTIONAL_SYNC_TEST_MOUNT_BLOCKDEV
FUNCTIONAL_SYNC_TEST_MOUNT_POINT
FUNCTIONAL_SYNC_TEST_LEN
FUNCTIONAL_SYNC_TEST_LOOP
```

Note: could be improved by checking for CONFIG_SYSVIPC in /proc/config.gz to verify that the required kernel features are present

MOUNT_BLOCKDEV and MOUNT_POINT are used by 'hd_test_mount_prepare' but are prefaced with FUNCTIONAL_SYNC_TEST or BENCHMARK_BONNIE

from clicking "Run Test", to executing code on the target... config.xml has the slave command: /home/jenkins/slave.jar -> which is a link to /home/jenkins/jta/engine/slave.jar

overlays.sh has "run_python \$OF_OVGEN ..." where OF_OVGEN is set to "\$JTA_SCRIPTS_PATH/ovgen/ovgen.py"

How is overlays.sh called?

it is sourced by /home/jenkins/scripts/benchmarks.sh and
/home/jenkins/scripts/functional.sh

functional.sh is sourced by each Functional.foo script.

For Functional.sync_test:

```
Functional.sync_test/config.xml
for the attribute <udson.tasks.Shell> (in <builders>)
    <command>....
        source $JTA_TESTS_PATH/$JOB_NAME/sync_test.sh</command>

sync_test.sh
    '. $JTA_SCRIPTS_PATH/functional.sh'
    'source $JTA_SCRIPTS_PATH/overlays.sh'
```

(continues on next page)

(continued from previous page)

```

'set_overlay_vars'
    (in overlays.sh)
    run_python $OF_OVGEN ($JTA_SCRIPTS_PATH/ovgen/ovgen.py) ...
        $OF_OUTPUT_FILE ($JTA_SCRIPTS_PATH/work/${NODE_NAME}_prolog.sh)
    generate xxx_prolog.sh
    SOURCE xxx_prolog.sh

functions.sh pre_test()

functions.sh build()
... test_build()

functions.sh deploy()

test_run()
    assert_define()
    functions.sh report()

```

17.2.3 NOTES about ovgen.py

What does this program do?

Here is a sample command line from a test console output:

```

python /home/jenkins/scripts/ovgen/ovgen.py \
  --classdir /home/jenkins/overlays//base \
  --ovfiles /home/jenkins/overlays//distribs/nologger.dist /home/jenkins/overlays//
↳boards/bbb.board \
  --testplan /home/jenkins/overlays//testplans/testplan_default.json \
  --specdir /home/jenkins/overlays//test_specs/ \
  --output /home/jenkins/work/bbb_prolog.sh

```

So, ovgen.py takes a classdir, a list of ovfiles a testplan and a specdir, and produces a xxx_prolog.sh file, which is then sourced by the main test script

Here is information about ovgen.py source:

```

Classes:
  OFClass
  OFLayer
  TestSpecs

```

```

Functions:
  parseOFVars - parse Overlay Framework variables and definitions
  parseVars - parse variables definitions
  parseFunctionBodyName
  parseFunction
  baseParseFunction
  parseBaseFile
  parseBaseDir
  parseInherit
  parseInclude

```

(continues on next page)

(continued from previous page)

```
parseLayerVarOverride
parseLayerFuncOverride
parseLayerVarDefinition
parseLayerCapList - look for BOARD.CAP_LIST
parseOverrideFile
generateProlog
generateSpec
parseGenTestPlan
parseSpec
parseSpecDir
run
```

17.2.4 Sample generated test script

bbb_prolog.sh is 195 lines, and has the following vars and functions:

```
from class:base-distrib:
    ov_get_firmware()
    ov_rootfs_kill()
    ov_rootfs_drop_caches()
    ov_rootfs_oom()
    ov_rootfs_sync()
    ov_rootfs_reboot()
    ov_rootfs_state()
    ov_logger()
    ov_rootfs_logread()

from class:base-board:
    LTP_OPEN_POSIX_SUBTEST_COUNT_POS
    MMC_DEV
    SRV_IP
    SATA_DEV
    ...
    JTA_HOME
    IPADDR
    PLATFORM=""
    LOGIN
    PASSWORD
    TRANSPORT
    ov_transport_cmd()
    ov_transport_put()
    ov_transport_get()

from class:base-params:
    DEVICE
    PATH
    SSH
    SCP

from class:base-funcs:
    default_target_route_setup()
```

(continues on next page)

(continued from previous page)

```

from testplan:default:
    BENCHMARK_DHRYSTONE_LOOPS
    BENCHMARK_<TESTNAME>_<VARNAME>
    ...
    FUNCTIONAL_<TESTNAME>_<VARNAME>

```

17.3 Logs

When a test is executed, several different kinds of logs are generated: devlog, systemlogs, the testlogs, and the console log.

17.3.1 created by Jenkins

- console log
 - this is located in `/var/lib/jenkins/jobs/<test_name>/builds/<build_id>/log`
 - it has the output from running the test script (on the host)

17.3.2 created by ftc

- console log
 - if 'ftc' was used to run the test, then the console log is created in the log directory, which is: `/fuego-rw/logs/<test_name>/<board>.<spec>.<build_id>.<build_number>/`
 - it is called `consolelog.txt`

17.3.3 created by the test script

- these are created in the directory: `/fuego-rw/logs/<test_name>/<board>.<spec>.<build_id>.<build_number>/`
- devlog has a list of commands run on the board during the test
 - named `devlog.txt`
- system logs have the log data from the board (e.g. `/var/log/messages`) before and after the test run:
 - named: `syslog.before.txt` and `syslog.after.txt`
- the test logs have the actual output from the test program on the target
 - this is completely dependent on what the test program outputs
 - named: `testlog.txt`
 - * this is the 'raw' log
 - there may be 'parsed' logs, which is the log filtered by `log_compare` operations:
 - * this is named: `testlog.p.txt` or `testlog.n.txt`
 - * the 'p' indicated positive results and the 'n' indicates negative results

17.4 Core scripts

The test script is sourced by the Fuego `main.sh` script

This script sources several other scripts, and ends up including `fuego_test.sh`

- load overlays and `set_overlay` vars
- `pre_test $TEST_DIR`
- build
- deploy
- test_run
- `set_testres_file`, `bench_processing`, `check_create_logrun` (if a benchmark)
- `get_testlog $TESTDIR`, `test_processing` (if a functional test)
- `get_testlog $TESTDIR` (if a stress test)
- `test_processing` (if a regular test)

functions available to test scripts: See Test Script APIs

Benchmark tests must provide a `parser.py` file, which extracts the benchmark results from the log data.

It does this by doing the following:

```
import common as plib
f = open(plib.TEST_LOG)
lines = f.readlines()
((parse the data))
```

This creates a dictionary with a key and value, where the key matches the string in the `reference.log` file

The `parser.py` program builds a dictionary of values by parsing the log from the test (basically the test output). It then sends the dictionary, and the pattern for matching the reference log test criteria to the routine: `common.py:process_data()`

It defines `ref_section_pat`, and passes that to `process_data()` Here are the different patterns for `ref_section_pat`:

```
9  "\[[\w]+.[gle]{2}\]"
1  "\[[w*.[gle]{2}\]"
1  "^\\[[d\\w_ .]+.[gle]{2}\\]"
1  "^\\[[d\\w_ -]+.[gle]{2}\\]"
1  "^\\[[w\\d&._/()]+.[gle]{2}\\]"
4  "^\\[[w\\d._]+.[gle]{2}\\]"
2  "^\\[[w\\d\\s_ -]+.[gle]{2}\\]"
3  "^\\[[w\\d_ ./]+.[gle]{2}\\]"
5  "^\\[[w\\d_ .]+.[gle]{2}\\]"
1  "^\\[[w\\d_\\- .]+.[gle]{2}\\]"
1  "^\\[[w]+.[gle]{2}\\]"
1  "^\\[[w_ .]+.[gle]{2}\\]"
```

Why are so many different ones needed?? Why couldn't the syntax be: `<var-name> <test> <value>` on one line?

It turns out this is processed by an 'awk' script. thus the weird syntax. We should get rid of the awk script and use python instead.

17.4.1 How is benchmarking graphing done?

See Benchmark parser note

17.4.2 docker tips

See Docker Tips

LICENSE AND CONTRIBUTION POLICY

18.1 License

Fuego has the following license policy.

Fuego consists of several parts, and includes source code from a number of different external test projects.

18.1.1 Default license

The default license for Fuego is the BSD 3-Clause license, as indicated in the LICENSE file at the top of the ‘fuego’ and ‘fuego-core’ source repositories.

If a file does not have an explicit license, or license indicator (such as SPDX identifier) in the file, then that file is covered by the default license for the project, with the exceptions noted below for “external test materials”.

When making contributions, if you do NOT indicate an alternative license for your contribution, the contribution will be assigned the license of the file to which the contribution applies (which may be the default license, if the file contains no existing license indicator).

Although we may allow for other licenses within the Fuego project in order to accommodate external software added to our system, our preference is to avoid the proliferation of licenses in the Fuego code itself.

External test materials

Individual tests in Fuego consist of files in the directory: `fuego-core/tests/<test_name>` (which is known as the test home directory), and may include two types of materials:

1. Fuego-specific files
2. files obtained from external sources, which have their own license.

The Fuego-specific materials consist of files such as: `fuego_test.sh`, `spec.json`, `criteria.json`, `test.yaml`, `chart_config.json`, and possibly others as created for use in the Fuego project. External test materials may consist of tar files, helper scripts and patches against the source in the tar files.

Unless otherwise indicated, the Fuego-specific materials are licensed under the Fuego default license, and the external test materials are licensed under their own individual project license - as indicated in the test source.

In some cases, there is no external source code, but only source that is originally written for Fuego and stored in the test home directory. This commonly includes tests based on a single shell script, that is written to be deployed to the Device Under Test by `fuego_test.sh`. Unless otherwise indicated, these files (source and scripts) are licensed under the Fuego default license.

If there is any ambiguity in the category of a particular file (external or Fuego-specific), please designate the intended license clearly in the file itself, when making a contribution.

18.1.2 Copyright statements

Copyrights for individual contributions should be added to individual files, when the contributions warrant copyright assignment. Some trivial fixes to existing code may not need to have copyright assignment, and thus not every change to a file needs to include a copyright notice for the contributor.

18.1.3 License tags

Our preference is to use SPDX license identifier, rather than a license notice, to indicate the license of any materials in Fuego. Such identifiers and notices are only desired if the materials are not contributed under the default Fuego license of “BSD-3-Clause”.

In a test.yaml, please indicate the license of the upstream test program. If there is no upstream test program (ie, the test is self-contained), please specify the license of the Fuego test definition itself.

Please see <https://spdx.org/licenses/> for a list of SPDX license tags

18.1.4 Contributor agreement

The Fuego project does not require a signed Contributor License Agreement for contribution to the project. Instead, we utilize the following Developer Certificate of Origin that was copied from the Linux kernel.

Each contribution to Fuego must be accompanied by a Signed-off-by line in the patch or commit description, which indicates agreement to the following:

By making a contribution to this project, I certify that:

- (a) The contribution was created **in** whole **or in** part by me **and** I have the right to submit it under the **open** source license indicated **in** the file; **or**
- (b) The contribution **is** based upon previous work that, to the best of my knowledge, **is** covered under an appropriate **open** source license **and** I have the right under that license to submit that work **with** modifications, whether created **in** whole **or in** part by me, under the same **open** source license (unless I am permitted to submit under a different license), **as** indicated **in** the file; **or**
- (c) The contribution was provided directly to me by some other person who certified (a), (b) **or** (c) **and** I have **not** modified it.
- (d) I understand **and** agree that this project **and** the contribution are public **and** that a record of the contribution (including **all** personal information I submit **with** it, including my sign-off) **is** maintained indefinitely **and** may be redistributed consistent **with** this project **or** the **open** source license(s) involved.

Note: Please note that an “official” DCO at the web site <https://developercertificate.org/> has additional text (an LF copyright, address, and statement of non-copyability). All of these extra items are either nonsense or problematical

in some legal sense. The above is a quote of a portion of the document found in the Linux kernel guide for submitting patches. See <https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/tree/Documentation/process/submitting-patches.rst> (copied in March, 2018).

Each commit must include a DCO which looks like this

`Signed-off-by: Joe Smith <joe.smith@email.com>`

The project requires that the name used is your real name. Neither anonymous contributors nor those utilizing pseudonyms will be accepted.

You may type this line on your own when writing your commit messages. However, Git makes it easy to add this line to your commit messages. Make sure the `user.name` and `user.email` are set in your git configs. Use `'-s'` or `'-signoff'` options to `'git commit'` to add the Signed-off-by line to the end of the commit message.

18.2 Submitting contributions

Please format contributions as a patch, and send the patch to the [Fuego mailing list](#)

Before making the patch, please verify that you have followed our preferred Coding style.

We follow the style of patches used by the Linux kernel, which is described here: <https://www.kernel.org/doc/html/latest/process/submitting-patches.html>

Not everything described there applies, but please do the following:

- Use a Signed-off-by line
- Send patch in plain text
- Include PATCH in the subject line
- Number patches in a series (1/n, 2/n, .. n/n)
- Use a subsystem prefix on the patch subject line
 - Patch subject should have: “subsystem: description”
 - In the case of modifications to a test, the subject should have: “test: description” (that is, the test is the subsystem name)
 - The test name can be the short name, if it is unambiguous
 - * That is, please don’t use the ‘Functional’ or ‘Benchmark’ prefix unless there are both types of tests with the same short name
- Describe your changes in the commit message body

18.2.1 Creating patches

If you use git, it’s easy to create a patch (or patch series), using `git format-patch`. Or, you can go directly to e-mailing a patch or patch series using `git send-email`.

18.2.2 Alternative submission method

I also allow patches as attachments to an e-mail to the list. This is something NOT supported by the Linux kernel community.

If the patch is too big (greater than 300K), then please add it to a public git repository, and let me know the URL for the repository. I can add a remote for the repo, and fetch it and cherry pick the patch. I prefer doing a fetch and cherry-pick to a pull request.

While I will sometimes process patches through a repo, it is strongly preferred for patches to go through the mailing list as plain text, so that community members can review the patch in public.

INTEGRATION WITH TTC

This page describes how to use Fuego with `ttc`. `ttc` is a tool used for manual and automated access to and manipulation of target boards. It is a tool developed by Tim Bird and used at Sony for managing their board farms, and for doing kernel development on multiple different target boards at a time (including especially boards with different processors and architectures.)

This page describes how `ttc` and Fuego are integrated, so that the Fuego test framework can use `ttc` as its transport mechanism.

You can find more information about `ttc` on the embedded Linux wiki at: http://elinux.org/Ttc_Program_Usage_Guide

19.1 Outline of supported functionality

Here is a rough outline of the support for `ttc` in Fuego:

- Integration for the tool and helper utilities in the container build
 - When the docker container is built, `ttc` is downloaded from Github and installed into the Docker image.
 - During this process, the path to the `ttc.conf` file is changed from `/etc/ttc.conf` to `/fuego-ro/conf/ttc.conf`
- `ttc` is a valid transport option in Fuego
 - You can specify `ttc` as the ‘transport’ for a board, instead of `ssh`
- `ttc` added support for `-r` as an option to the `ttc cp` command
 - This is required since Fuego uses `-r` extensively to do recursive directory copies (See `Transport_notes` for details)
- The Fuego core scripts have been modified to avoid using wildcards on `get` operations
- A new test called `Functional.fuego_transport` has been added
 - This tests use of wildcards, multiple files and directories and directory recursion with the Fuego `put` function.
 - It also indirectly tests the `get` function (the other major Fuego transport function), because logs are obtained during the test.

19.2 Supported operations

`ttc` has several sub-commands. Fuego currently only uses the following `ttc` sub-commands:

- `ttc run` - to run a command on the target
- `ttc cp` - to get a file from the target, and to put files to the target

Note that some other commands, such as `ttc reboot` are not used, in spite of there being similar functionality provided in fuego (see function `target reboot` and function `ov rootfs reboot`).

Finally, other commands, such as `ttc get_kernel`, `ttc get_config`, `ttc kbuild` and `ttc kinstall` are not used currently. These may be used in the future, when Fuego is expanded to have a focus on tests that require kernel rebuilding.

19.3 Location of `ttc.conf`

Normally, `ttc` on a host uses the default configuration file at `/etc/ttc.conf`. Fuego modifies the `ttc` installed inside the Fuego docker container, so that it uses the configuration file located at `/fuego-ro/conf/ttc.conf` as its default.

During Fuego installation, `/etc/ttc.conf` is copied to `/fuego-ro/conf` from the host machine, if it is present (and a copy of `ttc.conf` is not already there).

19.4 Steps to use `ttc` with a target board

Here is a list of steps to set up a target board to use `ttc`. These steps assume you have already added a board to fuego following the steps described in [Adding a board](#).

- If needed, create your Docker container using `docker-create-usb-privileged-container.sh`
 - This may be needed if you are using `ttc` with board controls that require access to USB devices (such as the Sony debug board)
 - Use the `--priv` option with `install.sh`, as documented in [Installing Fuego](#).
- Make sure that `fuego-ro/conf/ttc.conf` has the definitions required for your target board
 - Validate this by doing `ttc list` to see that the board is present, and `ttc run` and `ttc cp` commands, to test that these operations work with the board, from inside the container.
- Edit the fuego board file (found in `/fuego-ro/conf/boards /{board_name}.board`)
 - Set the `TRANSPORT` to “`ttc`”
 - Set the `TTC_TARGET` variable to the name for the target used by `ttc`
 - See the following example, for a definition for a target named ‘bbb’ (for my Beaglebone black board):

```
TRANSPORT=ttc
TTC_TARGET=bbb
```

19.5 modify your copy_to_cmd

In your `ttc.conf` file, you may need to make changes to the `copy_to_cmd` definitions for boards used by Fuego. Fuego allows programs to pass a `-r` argument to its internal `put` command, which in turn invokes `ttc`'s `cp` command, with the source as target and destination as the host. In other words, it ends up invoking `ttc`'s `copy_from_cmd` for the indicated target.

All instances of `copy_to_cmd` should be modified to reference a new environment variable `$copy_args`, and they should support the use of `-r` in the command arguments.

Basically, if a Fuego test uses `put -r` at any point, this needs to be supported by `ttc`. `ttc` will pass any `'-r'` seen to the subcommand in the environment variable `$copy_args`, where you can use it as needed with whatever sub-command (`cp`, `scp`, or something else) that you use to execute a `copy_to_cmd`.

See `ttc.conf.sample` and `ttc.conf.sample2` for usage examples.

WORKING WITH REMOTE BOARDS

Here are some general tips for working with remote boards (that is, boards in remote labs)

20.1 using a jump server

If you have an SSH jump server, then you can access machine directly in another lab, using the SSH ProxyCommand in the host settings for a board.

I found this page to be helpful: <https://www.tecmint.com/access-linux-server-using-a-jump-host/>

You should try to make each leg of the jump (from local machine to jump server, and from jump server to remote machine) password-less.

I found that if my local machine's public key was in the remote machine's authorized keys file, then I could log in without a password, even if the jump server's public key was not in the remote machine's authorized keys file.

20.2 Using ttc transport remotely

If you have a server that already has ttc configured for a bunch of board, you can accomplish a lot just by referencing ttc commands on that server.

For example, in your local ttc.conf, you can put:

```
PASSWORD=foo
USER=myuser
SSH_ARGS=-o UserKnownHostsFile=/dev/null -o StrictHostKeychecking=no -o LogLevel=QUIET

pos_cmd=ssh timdesk ttc %(target)s pos
off_cmd=ssh timdesk ttc %(target)s off
on_cmd=ssh timdesk ttc %(target)s on
reboot_cmd=ssh timdesk ttc %(target)s reboot

login_cmd=sshpass -p %(PASSWORD)s ssh %(SSH_ARGS)s -x %(USER)s@%(target)s
run_cmd=sshpass -p %(PASSWORD)s ssh %(SSH_ARGS)s -x %(USER)s@%(target)s "$COMMAND"
copy_to_cmd=sshpass -p %(PASSWORD)s scp %(SSH_ARGS)s $src %(USER)s@%(target)s:/$dest
copy_from_cmd=sshpass -p %(PASSWORD)s scp %(SSH_ARGS)s %(USER)s@%(target)s:/$src $dest
```

Note: Note that `ttc status <remote-board>` command does not work with ttc version 1.4.4. This is due to internal usage of `%(ip_addr)s` in the function `network_status()`, which will not be correct for the remote-board.

20.3 Setting up ssh ProxyCommand in the Fuego docker container

Please note that tests in Fuego are executed inside the Docker container as user 'jenkins'.

In order to set up password-less operation, or use of a jump server or ProxyCommand, you have to add appropriate items (config and keys) to `/var/lib/jenkins/.ssh`.

Please note that this may make your Docker container a security risk, as it may expose your private keys to tests. Please use caution when adding private keys or other sensitive security information to the Docker container.

API REFERENCE

This page documents the Fuego core API, which consists of a set of functions callable by a Fuego test, and a set of variables available to tests from the Fuego system.

The following functions are available to test scripts.

FIXTHIS - core functions: should document which of the functions are internal, and which are intended for test script use

FUNCTIONS AVAILABLE FOR TESTS TO CALL

These are the functions that can be called by a test's script (`fuego_test.sh`)

You could consider this the library of functions provided by Fuego for the test operation.

22.1 Functions for interacting with the target

These commands are used to provide generic methods to execute commands on the board and to copy files to and from the board. Fuego insulates the test itself from the specific commands needed to interact with the board (via, for example, ssh, serial console, adb, telnet, ttc, etc.)

- `cmd` - execute a command on the target
- `get` - get a file or files from the target
- `put` - put a file or files on the target
- `safe_cmd` - execute a command, handling out-of-memory conditions

22.2 Functions for checking dependencies and requirements

These function allow for checking that the test has required variables defined, and that the board has required programs, libraries and modules.

- ``assert_has_program`_` - check that a program is present on the target, and abort the test if it is missing
- ``check_has_program`_` - check to see if a specified program is on the target
- ``assert_has_module`_` - check to see if a specified kernel module is on the board
- ``assert_define`_` - to be used externally - check that an environment variable is defined
- ``get_program_path`_` - find path to execute a program on the target board
- ``is_abs_path`_` - check if a path is absolute (starts with /) or not
- ``is_empty`_` - fail if an environment variable is empty
- ``is_on_target`_` - check if a program or file is on the target board
- ``is_on_target_path`_` - check if a program is in one of the directories in the PATH on the target board

22.3 Functions for preparing board and executing tests

- ``function_hd_test_mount_prepare`_` - make sure test filesystems are mounted and ready for testing
- ``function_hd_test_clean_umount`_` - unmount test filesystems
- ``function_log_this`_` - used to execute operations on the host (as opposed to on the board), and log their output
- ``function_report`_` - execute command on board and put stdout into the test log
- ``function_report_append`_` - execute command on board and append stdout to log
- ``function_report_live`_` - execute command on board, and capture stdout to a log on the host

22.4 Functions for parsing results

This is a simple function for checking results in a test log

- ``log_compare`_` - check results of test by scanning the test log for a regular expression

22.5 Functions for cleanup

These functions may be used

- ``kill_procs`_` - kill processes on the board
- ``target_reboot`_` - reboot the board

22.6 For batch tests

- ``allocate_next_batch_id`_` - allocate and return the next unique batch id
- ``run_test`_` - run another test from the current test (almost always a batch test)

22.7 Misceleneous functions

- `run_python_` - used to run a python program on the host (inside the docker container)

22.8 For printing messages at various message output levels

- ``dprint`_` - print a 'debug' message
- ``vprint`_` - print a 'verbose' message
- ``iprint`_` - print an 'info' message
- ``wprint`_` - print a 'warning' message
- ``eprint`_` - print an 'error' message

22.9 fuego_test.sh functions

This is a list of functions that a test can provide to the Fuego system (in the test's `fuego_test.sh` script):

These functions correspond roughly to the test phases, which are normally these phases (in this order):

<code>pre_test</code> , <code>pre_check</code> , <code>build</code> , <code>deploy</code> , <code>makepkg</code> , <code>run</code> , <code>post_test</code> , <code>processing</code>
--

Here are the main `fuego_test.sh` functions:

- `test_pre_check_` - check that the board has needed dependencies and attributes
- `test_build_` - build the test software
- `test_deploy_` - put the test software on the board
- `test_run_` - execute the test software on the board
- `test_processing_` - parse test output for results

Most tests will have most of these functions. But any test can omit functions that are not needed. For example, if a test has no dependencies, does not have a binary program that needs to be compiled, or any script that needs to be deployed to the board, a test might omit the `test_pre_check`, `test_build`, and `test_deploy` functions, and only have the `test_run` and `test_processing` phases.

Here are functions that are allowed in `fuego_test.sh`, that can be used to override the normal Fuego operations. Most tests will not include these functions.

- `test_snapshot_` - get board status and info (customize the “board snapshot” feature)
- `test_fetch_results_` - get test results from the board (customize the fetch operation)
- `test_cleanup_` - clean up the board after the test (customize the cleanup operation)

PARSER MODULE API

The file `common.py` is the python module for performing benchmark log file processing, and results processing and aggregation.

It is used by the `parser.py` program from the test directory, to process the log after each test run. The data from a test run is processed to:

- Check numeric values for pass/fail result(by checking against a reference threshold values)
- Determine the overall result of the test, based on potentially complex results criteria
- Save the data for use in history and comparison charts

23.1 Parser API

The following are functions used during log processing, by a test's `parser.py` program.

- `parse_log()` - parse the data from a test log
 - This routine takes a regular expression, with one or more groups, and results a list of tuples for lines that matched the expression
 - The tuples consist of the strings from the matching line corresponding to the regex groups
- `process()` - process results from a test
 - This routine takes a dictionary of test results, and does 3 things:
 - * Formats them into the `run.json` file (run results file)
 - * Detects pass or fail by using the specified pass criteria
 - * Formats the data into charts (plots and tables)
- `split_output_per_testcase()`
 - Split testlog into chunks accessible from the Jenkins user interface (one per testcase)

In general, a parser module will normally call `parse_log()`, then take the resulting list of matching groups to construct a dictionary to pass to the `process()` routine.

If the log file format is amendable, the parser module may also call `split_output_per_testcase()` to generate a set of files from the testlog, that can be referenced from the charts generated by the charting module.

Please see [parser.py](#) for more details and examples of use of the API.

23.2 Deprecated API

Note: The following information is for historical purposes only. Although the API is still present in Fuego, these APIs are deprecated.

In Fuego version 1.1 and prior, the following functions were used. These are still available for backwards compatibility with tests written for these versions of Fuego.

- `parse()`
- `process_data()`

(see [*parser.py*](#) for invocation details)

23.2.1 `parse()`

- input:
 - `cur_search_pattern` - compiled re search pattern
- output:
 - list of regular expression matches for each line matching the specified pattern

This routine scans the current log file, using a regular expression. It returns an re match object for each line of the log file that matches the expression.

This list is used to populate a dictionary of metric/value pairs that can be passed to the `process_data` function.

23.2.2 `process_data`

This is the main routine of the module. It processes the list of metrics, and populates various output files for test.

- input:
 - `ref_section_pat` - regular expression used to read reference.log
 - `cur_dict` - dictionary of metric/value pairs
 - `m` - indicates the size of the plot. It should be one of: 's', 'm', 'l', 'xl'
 - * if 'm', 'l', or 'xl' are used, then a multiplot is created
 - `label` - label for the plot

This routine has the following outline:

- `write_report_results`
- read the reference thresholds
- check the values against the reference thresholds
- store the plot data to a file (`plot.data`)
- create the plot
- save the plot to an image file (`plot.png`)

23.3 Developer notes

23.3.1 functions in common.py

- hls - print a big warning or error message
- parse_log(regex_str) - specify a regular expression string to use to parse lines in the log
 - this is a helper function that returns a list of matches (with groups) that the parser.py can use to populate its dictionary of measurements
- parse(regex_compiled_object)
 - similar to parse_log, but it takes a compiled regular expression object, and returns a list of matches (with groups)
 - this is deprecated, but left to support legacy tests
- split_tguid()
- split_test_id()
- get_test_case()
- add_results()
- init_run_data()
- get_criterion()
- check_measure()
- decide_status()
- convert_reference_log_to_criteria()
- load_criteria()
- apply_criteria()
- create_default_ref()
- prepare_run_data()
- extract_test_case_ids()
- update_results_json()
- delete()
- save_run_json()
- process(results)
 - results is a dictionary with
 - * key=test_case_id (not including measure name)
 - for a functional test, the test_case_id is usually “default.<test_name>”
 - * value=list of measures (for a benchmark)
 - * or value=string (PASS|FAIL|SKIP) (for a functional test)
- process_data(ref_sections_pat, test_results, plot_type, label)

23.3.2 call trees

```
process_data(ref_section_pat, test_results, plot_type, label)
    process_data(measurements)
        prepare_run_data(results)
            run_data = (prepare non-results data structure)
            ref = read reference.json
                or ref = create_default_ref(results)
            init_run_data(run_data, ref)
                (put ref into run_data structure)
                (mark some items as SKIP)
            add_results(results, run_data)
                for each item in results dictionary:
                    (check for results type: list or str)
                    if list, add measure
                    if str, set status for test_case
            apply_criteria(run_data)
                load_criteria()
                    (load criteria.json)
                    or convert_reference_log_to_criteria()
                check_measure()
                    get_criterion()
                decide_status()
                    get_criterion()
            save_run_json(run_data)
            update_results_json()
            (return appropriate status)
```

23.3.3 miscellaneous notes

- create_default_ref_tim (for docker.hello-fail.Functional.hello_world)
 - ref={ 'test_sets': [{ 'test_cases': [{ 'measurements': [{ 'status': 'FAIL', 'name': 'Functional' }], 'name': 'default' }], 'name': 'default' }}}
- create_default_ref
 - ref={ 'test_sets': [{ 'test_cases': [{ 'status': 'FAIL', 'name': 'default' }], 'name': 'default' }}}

23.3.4 data format and tguid rules

The current API and the old parser API take different data and allow different test identifiers. This sections explains the difference:

Data format for benchmark test with new API

- measurements[test_case_id] = [{"name": measure_name, "measure": value}]

Data format for benchmark test with old API:

- in reference.log
 - if tguid is a single word, then use that word as the measure name and “default” as the test_case.
 - * e.g. for benchmark.arm, the reference.log has “short”. This becomes the fully-qualified tguid: arm.default.arm.short:

- test_name = arm, test_case = default, test_case_id = arm, measure = short

Data format for functional tests with new API and the old API is the same:

- e.g. measurements["status"] = "PASS|FAIL"

CORE INTERFACES

This page documents the interface between the Jenkins front end and the Fuego core engine. See also *Variables*

24.1 From Jenkins to Fuego

24.1.1 Environment variables passed during a build

Built-in Jenkins variables for shell script jobs

BUILD_ID

The current test run id. As of (at least) Jenkins version 2.32.1, this is the same as the BUILD_NUMBER.

BUILD_NUMBER

The current test run number, such as “153”. This appears to be selected as the next numerical number in sequence, by the Jenkins system, at the start of a job.

BUILD_TAG

String of “jenkins-\${JOB_NAME}-\${BUILD_NUMBER}”. Convenient to put into a resource file, a jar file, etc for easier identification.

BUILD_URL

Full URL of this test run, like <http://server:port/jenkins/job/foo/15/>

EXECUTOR_NUMBER

The unique number that identifies the current executor (among executors of the same machine) that’s carrying out this test run. This is the number you see in the “test executor status”, except that the number starts from 0, not 1.

JENKINS_HOME

The absolute path of the directory assigned on the master node for Jenkins to store data.

JENKINS_URL

Full URL of Jenkins, like <http://server:port/jenkins/>

JOB_NAME

Name of Jenkins job for this test. In Fuego, this will be something like: “myboard.default.Functional.foo” or “myboard.default.Benchmark.bar”. The job name has the form: {board}.{spec}.{type}.{test_name}

JOB_URL

Full URL of this test or test suite, like <http://server:port/jenkins/job/myboard.default.Functional.foo/>

NODE_LABELS

Whitespace-separated list of labels that are assigned to the node.

NODE_NAME

Name of the slave if the test run is on a slave, or “master” if run on master. In the case of Fuego, this is the board name (e.g. ‘beagleboneblack’)

WORKSPACE

The absolute path of the directory assigned to the test run as a workspace. For Fuego, this is always /fuego-rw/buildzone. Note that this comes from “custom workspace” setting in the job definition.

Fuego variables passed from Jenkins system configuration

The following variables are defined at the system level, and are passed by Jenkins in the environment of every job that is executed:

FUEGO_RO

The location where Fuego read-only data (configuration, boards and toolchains) are located. Currently /fuego-ro in the docker container.

FUEGO_RW

The location where Fuego read-write data is located (directories for builds, logs (run data), and other workspace and scratchpad areas). Currently /fuego-rw.

FUEGO_CORE

The location where the Fuego script system and tests are located. Currently /fuego-core.

Fuego variables passed from Jenkins job definition

These variables are defined in the job definition for a test:

Device

This is the target board to run the test on.

Reboot

Indicates to reboot the target device before running the test

Rebuild

Indicates that build instances of the test suite should be deleted, and the test suite rebuilt from the tarball

Target_PreCleanup

Indicates to clean up test materials left from a previous run of the test, before the test begins.

Target_PostCleanup

Indicates to clean up test materials after the test ends.

TESTPLAN

This has the name of the testplan used for this test job. Note that this selected by the end user from the output of getTestplans.groovy. (example value: testplan_default)

TESTNAME

Has the base name of the test (e.g. LTP vs Functional.LTP)

TESTSPEC

This has the name of the spec used for this test job

FUEGO_DEBUG

Can have a 1 to indicate verbose shell script output

24.2 From Fuego to Fuego

DISTRIB

Indicates the distribution file for the board. This comes from the board file. It's primary purpose is to select the logging features of the distribution on the target (to indicate whether there's a logger present on target). The value is often 'distrib/nologread.dist'

OF_BOARD_FILE

Full path to the board file

OF_CLASSDIR

full path to the overlay class directory (usually /home/jenkins/overlays//base)

OF_CLASSDIR_ARGS

argument specifying the overlay class directory (usually '-classdir /home/jenkins/overlays//base')

OF_DEFAULT_SPECDIR

path to directory containing test specs (usually /home/jenkins/overlays//test_specs)

OF_DISTRIB_FILE

path to the distribution file for the target board (often /home/jenkins/overlays//distrib/nologger.dist)

OF_OVFILES

FIXTHIS - document what OF_OVFILES is for

OF_OVFILES_ARGS

FIXTHIS - document what OF_OVFILES_ARGS is for

OF_ROOT

root directory for overlay generator (usually /home/jenkins/overlays/)

OF_SPECDIR_ARGS

argument to specify the test spec directory (usually '-specdir /home/jenkins/overlays//test_specs/')

OF_TESTPLAN_ARGS

OF_TESTPLAN

full path to the JSON test plan file for this test (often /home/jenkins/overlays//testplans/testplan_default.json)

OF_TESTPLAN_ARGS

argument specifying the path to the testplan (often '-testplan /home/jenkins/overlays//testplans/testplan_default.json')

TEST_HOME

home directory for the test materials for this test (example: /home/jenkins/tests/Functional.bc)

TESTDIR

base directory name of the test (example: Functional.bc)

TRIPLET

FIXTHIS - document TRIPLET

24.2.1 Deprecated

The following variables are no longer used in Fuego:

FUEGO_ENGINE_PATH

(deprecated in Fuego 1.1 - use '\$FUEGO_CORE/engine' now)

FUEGO_PARSER_PATH

(deprecated in Fuego 1.1)

24.3 Example Values

Here are the values from a run using the Jenkins front-end with job bbb.default.Functional.hello_world:

(these are sorted alphabetically):

```
AR=arm-linux-gnueabihf-ar
ARCH=arm
AS=arm-linux-gnueabihf-as
BUILD_DISPLAY_NAME=#2
BUILD_ID=2
BUILD_NUMBER=2
BUILD_TAG=jenkins-bbb.default.Functional.hello_world-2
BUILD_TIMESTAMP=2017-04-10_21-55-26
CC=arm-linux-gnueabihf-gcc
CONFIGURE_FLAGS=--target=arm-linux-gnueabihf --host=arm-linux-gnueabihf --build=x86_64-
↳ unknown-linux-gnu
CPP=arm-linux-gnueabihf-gcc -E
CROSS_COMPILE=arm-linux-gnueabihf-
CXX=arm-linux-gnueabihf-g++
CXXCPP=arm-linux-gnueabihf-g++ -E
EXECUTOR_NUMBER=0
FUEGO_CORE=/fuego-core
FUEGO_RO=/fuego-ro
FUEGO_RW=/fuego-rw
FUEGO_START_TIME=1491861326786
HOME=/var/lib/jenkins
HOST=arm-linux-gnueabihf
HUDSON_COOKIE=1b9620a3-d550-4cb1-afb1-9c5a29650c14
HUDSON_HOME=/var/lib/jenkins
HUDSON_SERVER_COOKIE=2334aa4d37eae7a4
JENKINS_HOME=/var/lib/jenkins
JENKINS_SERVER_COOKIE=2334aa4d37eae7a4
JOB_BASE_NAME=bbb.default.Functional.hello_world
JOB_DISPLAY_URL=http://unconfigured-jenkins-location/job/bbb.default.Functional.hello_
↳ world/display/redirect
JOB_NAME=bbb.default.Functional.hello_world
LD=arm-linux-gnueabihf-ld
LDFLAGS=---sysroot / -lm
LOGDIR=/fuego-rw/logs/Functional.hello_world/bbb.default.2.2
LOGNAME=jenkins
MAIL=/var/mail/jenkins
NODE_LABELS=bbb
```

(continues on next page)

(continued from previous page)

```

NODE_NAME=bbb
PATH=/usr/local/bin:/usr/local/bin:/usr/bin:/bin:/usr/local/games:/usr/games
PREFIX=arm-linux-gnueabihf
PWD=/fuego-rw/buildzone
RANLIB=arm-linux-gnueabihf-ranlib
Reboot=false
Rebuild=true
RUN_CHANGES_DISPLAY_URL=http://unconfigured-jenkins-location/job/bbb.default.Functional.
↳hello_world/2/display/redirect?page=changes
RUN_DISPLAY_URL=http://unconfigured-jenkins-location/job/bbb.default.Functional.hello_
↳world/2/display/redirect
SDKROOT=/
SHELL=/bin/bash
SHLVL=3
Target_PostCleanup=true
Target_PreCleanup=true
TERM=xterm
TESTDIR=Functional.hello_world
TESTNAME=hello_world
TESTSPEC=default
USER=jenkins
WORKSPACE=/fuego-rw/buildzone

```

24.4 From Fuego to Jenkins

This section describes some of the operations that Fuego core scripts (or a test) can perform to invoke an action by Jenkins during a test. To perform a Jenkins action, Fuego uses Jenkins' REST API using the `wget` command.

- To abort a job, fuego does:
 - `wget -qO- ${BUILD_URL}/stop`
 - This is called by `common.sh: abort_job()`
- To check if another test instance is running (do a lock check), fuego does:
 - `wget -qO- "$(cat ${LOCKFILE})/api/xml?xpath=*/building/text%28%29"`
 - * `LOCKFILE` was previously set to hold the contents: `${BUILD_URL}`, so this resolves to:
 - `wget -qO- ${BUILD_URL}/api/xml?xpath=*/building/text()`
 - This is called by `functions.sh:concurrent_check()`

24.4.1 Jenkins python module

Fuego's `ftc` command uses the 'jenkins' python module to perform a number of operations with Jenkins. This module is used to:

- list nodes
- add nodes
- remove nodes
- list jobs
- build jobs
- remove jobs
- re-synch build data for a job (using `get_job_config()` and `reconfig_job()`)
- add view

Note: Fuego uses `jenkins-cli` to add jobs (described next).

24.4.2 Jenkins-cli interface

You can run Jenkins commands from the command line, using the pre-installed `jenkins-cli` interface. This is used by Fuego's `ftc` command to create jobs.

`jenkins-cli.jar` is located in the Docker container at:

```
/var/cache/jenkins/war/WEB-INF/jenkins-cli.jar
```

See <https://wiki.jenkins-ci.org/display/JENKINS/Jenkins+CLI> for information about using this plugin.

Here is a list of available commands for this plugin:

```
build
  Runs a test, and optionally waits until its completion.
cancel-quiet-down
  Cancel the effect of the "quiet-down" command.
clear-queue
  Clears the test run queue
connect-node
  Reconnect to a node
console
  Retrieves console output of a build
copy-job
  Copies a test.
create-job
  Creates a new test by reading stdin as a configuration XML file.
delete-builds
  Deletes test record(s).
delete-job
  Deletes a test
delete-node
  Deletes a node
```

(continues on next page)

(continued from previous page)

```

disable-job
    Disables a test
disconnect-node
    Disconnects from a node
enable-job
    Enables a test
get-job
    Dumps the test definition XML to stdout
groovy
    Executes the specified Groovy script.
groovysh
    Runs an interactive groovy shell.
help
    Lists all the available commands.
install-plugin
    Installs a plugin either from a file, an URL, or from update center.
install-tool
    Performs automatic tool installation, and print its location to
    stdout. Can be only called from inside a test run.
keep-build
    Mark the test run to keep the test run forever.
list-changes
    Dumps the changelog for the specified test(s).
list-jobs
    Lists all tests in a specific view or item group.
list-plugins
    Outputs a list of installed plugins.
login
    Saves the current credential to allow future commands to run
    without explicit credential information.
logout
    Deletes the credential stored with the login command.
mail
    Reads stdin and sends that out as an e-mail.
offline-node
    Stop using a node for performing test runs temporarily, until the
    next "online-node" command.
online-node
    Resume using a node for performing test runs, to cancel out the
    earlier "offline-node" command.
quiet-down
    Quiet down Jenkins, in preparation for a restart. Don't start
    any test runs.
reload-configuration
    Discard all the loaded data in memory and reload everything from
    file system. Useful when you modified config files directly on disk.
restart
    Restart Jenkins
safe-restart
    Safely restart Jenkins
safe-shutdown
    Puts Jenkins into the quiet mode, wait for existing test runs to

```

(continues on next page)

(continued from previous page)

```
be completed, and then shut down Jenkins.
session-id
  Outputs the session ID, which changes every time Jenkins restarts
set-build-description
  Sets the description of a test run.
set-build-display-name
  Sets the displayName of a test run
set-build-result
  Sets the result of the current test run. Works only if invoked
  from within a test run.
shutdown
  Immediately shuts down Jenkins server
update-job
  Updates the test definition XML from stdin.
  The opposite of the get-job command
version
  Outputs the current version.
wait-node-offline
  Wait for a node to become offline
wait-node-online
  Wait for a node to become online
who-am-i
  Reports your credential and permissions
```

24.4.3 Scripts to process Fuego data

Benchmark parsing

In Fuego, Benchmark log parsing is done by a python system consisting of `parser.py` (from each test), `dataload.py` and utility functions in `fuego-core/engine/scripts/parser`

See Benchmark parser notes, [*parser.py*](#), [*reference.log*](#) and [*Parser module API*](#).

Postbuild action

In Fuego, Jenkins jobs are configured to perform a postbuild action, to set the description of a test with links to the test log (and possibly plot and other files generated in post-processing)

ADDING A NEW TEST

25.1 Overview of Steps

To add a new test to Fuego, you need to perform the following steps:

- 1. Decide on a test name and type
- 2. Make the test directory
- 3. Get the source (or binary) for the test
- 4. Write a test script for the test
- 5. Add the test_specs (if any) for the test
- 6. Add log processing to the test
- 6-a. (if a benchmark) Add parser.py and criteria and reference files
- 7. Create the Jenkins test configuration for the test

25.2 Decide on a test name

The first step to creating a test is deciding the test name. There are two types of tests supported by Fuego: functional tests and benchmark tests. A functional test either passes or fails, while a benchmark test produces one or more numbers representing some performance measurements for the system.

Usually, the name of the test will be a combination of the test type and a name to identify the test itself. Here are some examples: *bonnie* is a popular disk performance test. The name of this test in the fuego system is *Benchmark.bonnie*. A test which runs portions of the posix test suite is a functional test (it either passes or fails), and in Fuego is named *Functional.posixtestsuite*. The test name should be all one word (no spaces).

This name is used as the directory name where the test materials will live in the Fuego system.

25.3 Create the directory for the test

The main test directory is located in `/fuego-core/tests/<test_name>`

So if you just created a new Functional test called ‘foo’, you would create the directory:

- `/fuego-core/tests/Functional.foo`

25.4 Get the source for a test

The actual creation of the test program itself is outside the scope of Fuego. Fuego is intended to execute an existing test program, for which source code or a script already exists.

This page describes how to integrate such a test program into the Fuego test system.

A test program in Fuego is provided in source form so that it can be compiled for whatever processor architecture is used by the target under test. This source may be in the form of a tarfile, or a reference to a git repository, and one or more patches.

Create a tarfile for the test, by downloading the test source manually, and creating the tarfile. Or, note the reference for the git repository for the test source.

25.4.1 tarball source

If you are using source in the form of a tarfile, you add the name of the tarfile (called ‘tarball’) to the test script.

The tarfile may be compressed. Supported compression schemes, and their associated extensions are:

- uncompressed (extension=‘.tar’)
- compressed with gzip (extension=‘.tar.gz’ or ‘.tgz’)
- compressed with bzip2 (extension=‘.bz2’)

For example, if the source for your test was in the tarfile ‘foo-1.2.tgz’ you would add the following line to your test script, to reference this source:

```
tarball=foo-1.2.tgz
```

25.4.2 git source

If you are using source from an online git repository, you reference this source by adding the variables ‘gitrepo’ and ‘gitref’ to the test script.

In this case, the ‘gitrepo’ is the URL used to access the source, and the ‘gitref’ refers to a commit id (hash, tag, version, etc.) that refers to a particular version of the code.

For example, if your test program is built from source in an online ‘foo’ repository, and you want to use version 1.2 of that (which is tagged in the repository as ‘v1.2’, on the master branch, you might have some lines like the following in the test’s script.

```
gitrepo=http://github.com/sampleuser/foo.git
gitref=master/v1.2
```

25.4.3 script-based source

Some tests are simple enough to be implemented as a single script (that runs on the board). For these tests, no additional source is necessary, and the script can just be placed directly in the test's home directory. In *fuego_test.sh* you must set the following variable:

```
local_source=1
```

During the deploy phase, the script is sent to the board directly from the test home directory instead of from the test build directory.

25.5 Test script

The test script is a small shell script called *fuego_test.sh*. It specifies the source tarfile containing the test program, and provides implementations for the functions needed to build, deploy, execute, and evaluate the results from the test program.

The test script for a functional test should contain the following:

- source reference (either tarball or gitrepo and gitref)
- function *test_pre_check* (optional)
- function *test_build*
- function *test_deploy*
- function *test_run*
- function *test_processing*

The *test_pre_check* function is optional, and is used to check that the test environment and target configuration and setup are correct in order to run the test.

25.5.1 Sample test script

Here is the *fuego_test.sh* script for the test *Functional.hello_world*. This script demonstrates a lot of the core elements of a test script.:

```
#!/bin/bash

tarball=hello-test-1.0.tgz

function test_build {
    make
}

function test_deploy {
    put hello $BOARD_TESTDIR/fuego.$TESTDIR/
}

function test_run {
    report "cd $BOARD_TESTDIR/fuego.$TESTDIR; \
    ./hello $FUNCTIONAL_HELLO_WORLD_ARG"
}
```

(continues on next page)

(continued from previous page)

```
function test_processing {  
    log_compare "$TESTDIR" "1" "SUCCESS" "p"  
}
```

25.5.2 Description of test functions

The test functions (`test_build`, `test_deploy`, `test_run`, and `test_processing`) are fairly simple. Each one contains a few statements to accomplish that phase of the test execution.

You can find more information about each of these functions at the following links:

- `test_pre_check`
- `test_build`
- `test_deploy`
- `test_run`
- `test_processing`

25.6 Test spec

Another element of every test is the *test spec*. A file is used to define a set of parameters that are used to customize the test for a particular use case. Each “spec” defines a variant of the test, that can be executed.

You must define the test spec(s) for this test, and add an entry to the appropriate testplan for it.

Each test in the system must have a test spec file. This file is used to list customizable variables for the test.

If a test program has no customizable variables, or none are desired, then at a minimum a *default* test spec must be defined, with no test variables.

The test spec file is:

- named ‘spec.json’ in the test directory,
- in JSON format,
- provides a `testName` attribute, and a `specs` attribute, which is a list,
- may include any named spec you want, but must define at least the ‘default’ spec for the test
 - Note that the ‘default’ spec can be empty, if desired.

Here is an example one that defines no variables.

```
{  
    "testName": "Benchmark.OpenSSL",  
    "specs": {  
        "default": {}  
    }  
}
```

And here is the spec.json of the `Functional.hello_world` example, which defines three specs:

```
{
    "testName": "Functional.hello_world",
    "specs": {
        "hello-fail": {
            "ARG": "-f"
        },
        "hello-random": {
            "ARG": "-r"
        },
        "default": {
            "ARG": ""
        }
    }
}
```

25.7 Test results parser

Each test should also provide some mechanism to parse the results from the test program, and determine the success of the test.

For a simple Functional test, you can use the `log_compare` function to specify a pattern to search for in the test log, and the number of times that pattern should be found in order to indicate success of the test. This is done from the `test_processing` function in the test script.

Here is an example of a call to `log_compare`:

```
function test_processing {
    log_compare "$TESTDIR" "11" "^TEST.*OK" "p"
}
```

This example looks for the pattern `^TEST.*OK`, which finds lines in the test log that start with the word ‘TEST’ and are followed by the string ‘OK’ on the same line. It looks for this pattern 11 times.

`log_compare` can be used to parse the logs of simple tests with line-oriented output.

For tests with more complex output, and for Benchmark tests that produce numerical results, you must add a python program called ‘`parser.py`’, which scans the test log and produces a data structure used by other parts of the Fuego system.

See [`parser.py`](#) for information about this program.

25.8 Pass criteria and reference info

You should also provide information to Fuego to indicate how to evaluate the ultimate resolution of the test.

For a Functional test, it is usually the case that the whole test passes only if all individual test cases in the test pass. That is, one error in a test case indicates overall test failure. However, for Benchmark tests, the evaluation of the results is more complicated. It is required to specify what numbers constitute success vs. failure for the test.

Also, for very complicated Functional tests, there may be complicated results, where, for example, some results should be ignored.

You can specify the criteria used to evaluate the test results, by creating a ‘`criteria.json`’ file for the test.

Finally, you may wish to add a file that indicates certain information about the test results. This information is placed in the ‘reference.json’ file for a test.

Please see the links for those files to learn more about what they are and how to write them, and customize them for your system.

25.9 Jenkins job definition file

The last step in creating the test is to create the Jenkins job for it.

A Jenkins job describes to Jenkins what board to run the test on, what variables to pass to the test (including the test spec (or variant), and what script to run for the test.

Jenkins jobs are created using the command-line tool ‘ftc’.

A Jenkins job has the name `<node_name>.<spec>.<test_type>.<test_name>`

You create a Jenkins job using a command like the following:

- `$ ftc add-jobs -b myboard -t Functional.mytest [-s default]`

The ftc ‘add-jobs’ sub-command uses ‘-b’ to specify the board, ‘-t’ to specify the test, and ‘-s’ to specify the test spec that will be used for this Jenkins job.

In this case, the name of the Jenkins job that would be created would be:

- `myboard.default.Functional.mytest`

This results in the creation of a file called `config.xml`, in the `/var/lib/jenkins/jobs/<job_name>` directory.

25.10 Publishing the test

Tests that are of general interest should be submitted for inclusion into fuego-core.

Right now, the method of doing this is to create a commit and send that commit to the Fuego mailing list, for review, and hopefully acceptance and integration by the fuego maintainers.

In the future, a server will be provided where test developers can share tests that they have created in a kind of “test marketplace”. Tests will be available for browsing and downloading, with results from other developers available to compare with your own results. There is already preliminary support for packaging a test using the ‘ftc package-test’ feature. More information about this service will be made available in the future.

25.11 Technical Details

This section has technical details about a test.

25.11.1 Directory structure

The directory structure used by Fuego is documented at [[Fuego directories]]

25.11.2 Files

A test consists of the following files or items:

File or item	format	location	description	test type
con-fig.xml	Jenkins XML	/var/lib/jenkins/jobs/{test_name}	Has the Jenkins (front-end) configuration for the test	all
tarfile	tar format	/fuego-core/tests/{test_name}	Has the source code for the test program	all
patches	patch format	/fuego-core/tests/{test_name}	Zero or more patches to customize the test program (applied during the unpack phase	all
base script	shell script	/fuego-core/tests/{test_name} /fuego_test.sh	Is the shell script that implements the different test phases in Fuego	all
test spec	JSON	/fuego-core/tests/{test_name} /spec.json	Has groups of variables (and their values) that can be used with this test	all
test plan(s)	JSON	/fuego-core/overlays/testplan	Has the testplans for the entire system	all
parser	python	/fuego-core/tests/{test_name} /parser.py	Python program to parse benchmark metrics out of the log, and provide a dictionary to the Fuego plotter	all, but can be missing for functional tests
pass criteria	JSON	/fuego-core/tests/{test_name} /criteria.json	Has the “pass” criteria for the test	all
reference info	JSON	/fuego-core/tests/{test_name} /reference.json	Has additional information about test results, such as the units for benchmark measurements	benchmark only
reference.log	Fuego-specific	/fuego-core/tests/{test_name} /reference.log	Has the threshold values and comparison operators for benchmark metrics measured by the test	benchmark only
p/n logs	text	(deprecated)	Are logs with the results (positive or negative) parsed out, for determination of test pass/fail	functional only

USING BATCH TESTS

A “batch test” in Fuego is a Fuego test that runs a series of other tests as a group. The results of the individual tests are consolidated into a list of testcase results for the batch test.

Prior to Fuego version 1.5, there was a different feature, called “testplans”, which allowed users to compose sets of tests into logical groups, and run them together. The batch test system, introduced in Fuego version 1.5 replaces the testplan system.

26.1 How to make a batch test

A batch test consists of a Fuego test that runs other tests. A Fuego batch test is similar to other Fuego tests, in that the test definition lives in `fuego-core/tests/<test-name>`, and it consists of a `fuego_test.sh` file, a spec file, a `parser.py`, a `test.yaml` file and possibly other files.

The difference is that a Fuego batch test runs other Fuego tests, as a group. The batch test has a few elements that are different from other tests.

Inside the `fuego_test.sh` file, a batch test must define two main elements:

1. The testplan element
2. The `test_run` function, with commands to run other tests

26.1.1 Testplan element

The testplan element consists of data assigned to the shell variable `BATCH_TESTPLAN`. This variable contains lines that specify, in machine-readable form, the tests that are part of the batch job. The testplan is specified in json format, and is used to specify the attributes (such as timeout, flags, and specs) for each test. The testplan element is used by `ftc add-jobs` to create Jenkins jobs for each sub-test that is executed by this batch test.

The `BATCH_TESTPLAN` variable must be defined in the `fuego_test.sh` file. The definition must begin with a line starting with the string `'BATCH_TESTPLAN='` and end with a line starting with the string `'END_TESTPLAN'`. By convention this is defined as a shell “here document”, like this example:

```
BATCH_TESTPLAN=$(cat <<END_TESTPLAN
{
    "testPlanName": "foo_plan",
    "tests": [
        { "testName": "Functional.foo" },
        { "testName": "Functional.bar" }
    ]
}
```

(continues on next page)

(continued from previous page)

```
END_TESTPLAN
)
```

The lines of the testplan follow the format described at [Testplan_Reference](#). Please see that page for details about the plan fields and structure (the schema for the testplan data).

26.1.2 test_run function

The other element in a batch test's `fuego_test.sh` is a `test_run` function. This function is used to actually execute the tests in the batch.

There are two functions that are available to help with this:

- `allocate_next_batch_id`
- `run_test`

The body of the `test_run` function for a batch test usually has a few common elements:

- setting of the `FUEGO_BATCH_ID`
- execution of the sub-tests, using a call to the function `run_test` for each one

Here are the commands in the `test_run` function for the test `Functional.batch_hello`:

```
function test_run {
    export TC_NUM=1
    DEFAULT_TIMEOUT=3m
    export FUEGO_BATCH_ID="hello-
$(allocate_next_batch_id)"

    # don't stop on test errors
    set +e
    log_this "echo \"batch_id=$FUEGO_BATCH_ID\""
    TC_NAME="default"
    run_test Functional.hello_world
    TC_NAME="fail"
    run_test Functional.hello_world -s hello-fail
    TC_NAME="random"
    run_test Functional.hello_world -s hello-random
    set -e
}
```

Setting the batch_id

Fuego uses a `batch_id` to indicate that a group of test runs are related. Since a single Fuego test can be run in many different ways (e.g. from the command line or from Jenkins, triggered manually or automatically, or as part of one batch test or another), it is helpful for the run data for a test to be assigned a `batch_id` that can be used to generate reports or visualize data for the group of tests that are part of the batch.

A batch test should set the `FUEGO_BATCH_ID` for the run to a unique string for that run of the batch test. Each sub-test will store the batch id in its `run.json` file, and this can be used to filter run data in subsequent test operations. The Fuego system can provide a unique number, via the routine `allocate_next_batch_id`. By convention, the `batch_id` for a test is created by combining a test-specific prefix string with the number returned from `allocate_next_batch_id`.

In the example above, the prefix used is 'hello-', and this would be followed by a number returned by the function `allocate_next_batch_id`.

Executing sub-tests

The `run_test` function is used to execute the sub-tests that are part of the batch. The other portions of the example above show setting various shell variables that are used by `run_test`, and turning off 'errexit' mode while the sub-tests are running.

In the example above, `TC_NUM`, `TC_NAME`, and `DEFAULT_TIMEOUT` are used for various effects. These variables are optional, and in most cases a batch test can be written without having to set them. Fuego will generate automatic strings or values for these variables if they are not defined by the batch test.

Please see the documentation for `run_test` for details about the environment and arguments used when calling the function.

Avoiding stopping on errors

The example above shows use of `set +e` and `set -e` to control the shell's 'errexit' mode. By default, Fuego runs tests with the shell errexit mode enabled. However, a batch test should anticipate that some of its sub-tests might fail. If you want all of the tests in the batch to run, even if some of them fail, they you should use `set +e` to disable errexit mode, and `set -e` to re-enable it when you are done.

Of course, if you want the batch test to stop if one of the sub-tests fails, they you should control the errexit mode accordingly (for example, leaving it set during all sub-test executions, or disabling it or enabling it only during the execution of particular sub-tests).

Whether to manipulate the shell errexit mode or not depends on what the batch test is doing. If it is implementing a sequence of dependent test stages, the errexit mode should be left enabled. If a batch test is implementing a series of unrelated, independent tests, the errexit mode should be disabled and re-enabled as shown.

26.2 Test output

The `run_test` function logs test results in a format similar to TAP13. This consists of the test output, followed by a line starting with the batch id (inside double brackets), then "ok" or "not ok" to indicate the sub-test result, followed by the testcase number and testcase name.

A standard `parser.py` for this syntax is available and used by other batch tests in the system (See `fuego-core/tests/Functional.batch_hello/parser.py`)

26.3 Preparing the system for a batch job

In order to run a batch test from Jenkins, you need to define a Jenkins job for the batch test, and jobs for all of the sub-tests that are called by the batch test.

You can use `ftc add-jobs` with the batch test, and Fuego will create the job for the batch test itself as well as jobs for all of its sub-tests.

It is possible to run a batch test from the command line using 'ftc run-test', without creating Jenkins jobs. However if you want to see the results of the test in the Jenkins interface, then the Jenkins test jobs need to be defined prior to running the batch test from the command line.

26.4 Executing a batch test

A batch test is executed the same way as any other Fuego test. Once installed as a Jenkins job, you can execute it using the Jenkins interface (manually), or use Jenkins features to cause it to trigger automatically. Or, you can run the test from the command line using `ftc run-test`.

26.5 Viewing batch test results

You can view results from a batch test in two ways:

1. Inside the Jenkins interface, or
2. Using `ftc` to generate a report.

26.5.1 Jenkins batch test results tables

Inside the Jenkins interface, a batch job will display the list of sub-tests, and the PASS/FAIL status of each one. In addition, if there is a Jenkins job associated with a particular sub-test, there will be a link in the table cell for that test run, that you can click to see that individual test's result and data in the Jenkins interface.

26.5.2 Generating a report

You can view a report for a batch test, by specifying the `batch_id` with the `ftc gen-report` command.

To determine the `batch_id`, look at the log for the batch test (`testlog.txt` file). Or, generate a report listing the `batch_ids` for the batch test, like so:

```
$ ftc gen-report --where test=batch_<name> --fields timestamp,batch_id
```

Select an appropriate `batch_id` from the list that appears, and note it for use in the next command.

Now, to see the results from the individual sub-tests in the batch, use the desired `batch_id` as part of a “where” clause, like so:

```
$ ftc gen-report --where batch_id=<batch_id>
```

You should see a report with all sub-test results for the batch.

26.6 Miscellaneous notes

26.6.1 Timeouts

The timeout for a batch test should be long enough for all sub-tests to complete. When a batch test is launched from Jenkins, the board on which it will run is reserved and will be unavailable for tests until the entire batch is complete. Keep this in mind when executing batch tests that call sub-tests that have a long duration.

The timeout for individual sub-tests can be specified multiple ways. First, the timeout listed in the testplan (embedded in `fuego_test.sh` as the `BATCH_TESTPLAN` variable) is the one assigned to the Jenkins job for the sub-test, when jobs are created during test installation into Jenkins. These take effect when a sub-test is run independently from the batch test.

If you want to specify a non-default timeout for a test, then you must use a `--timeout` argument to the `run_test` function, for that sub-test.

FUEGO NAMING RULES

To add boards or write tests for Fuego, you have to create a number of files and define a number of variables. Here are some rules and conventions for naming things in Fuego:

27.1 Fuego test name

- a Fuego test name must have one of the following prefixes:
 - ‘Functional.’
 - ‘Benchmark.’
- the name following the prefix is known as the base test name, and has the following rules:
 - it may only use letters, numbers and underscores
 - * that is - no dashes
 - it may use upper and lower case letters
- All test definition materials reside in a directory with the full test name:
 - e.g. Functional.hello_world

27.2 Test files

- the base script file has the name **fuego_test.sh**
- the default spec file for a test has name **spec.json**
- the default criteria file for a test has the name **criteria.json**
- the reference file for a test has the name **reference.json**
- the parser module for a test has the filename **parser.py**

27.3 Test spec names

Test spec names are declared in the spec file for a test. These names should consist of letters, numbers and underscores only. That is, no periods should be used in spec names.

Every test should have at least one spec, called the ‘default’ spec. If no spec file exists for a test, then Fuego generates a ‘default’ spec (on-the-fly), which is empty (ie. has no test variables).

27.4 Board names

- boards are defined by files in the /fuego-ro/boards directory
- their filenames consists of the board name, with the suffix “.board”
 - e.g. beaglebone.board
- a board name should have only letters, numbers and underscores.
 - specifically, no dashes, periods, or other punctuation is allowed

27.5 Jenkins element names

Several of the items in Fuego are represented in the Jenkins interface. The following sections describe the names used for these elements.

27.5.1 Node name

- A Jenkins node corresponding to a board must have the same name as the board.
 - e.g. beaglebone

27.5.2 Job name

- A Jenkins job is used to execute a test.
- Jenkins job names should consist of these parts: <board>.<spec>.<test_name>
 - e.g. beaglebone.default.Functional.hello_world

27.6 Run identifier

A Fuego run identifier is used to refer to a “run” of a test - that is a particular invocation of a test and it’s resulting output, logs and artifacts. A run identifier should be unique throughout the world, as these are used in servers where data from runs from different hosts are stored.

The parts of a run id are separated by dashes, except that the separator between the host and the board is a colon.

A fully qualified (global) run identifier consist of the following parts:

- test name
- spec name

- build number
- the word *on*
- host
- board

FIXTHIS - global run ids should include timestamps to make them globally unique for all time

Example: Functional.LTP-quickhit-3-on-timdesk:beaglebone

A shortened run identifier may omit the *on* and *host*. This is referred to as a local run id, and is only valid on the host where the run was produced.

Example:

- Functional.netperf-default-2-minnow

27.7 timestamp

- A Fuego timestamp has the format: YYYY-mm-dd_HH:MM:SS
 - e.g. 2017-03-29_10:25:14
- times are expressed in localtime (relative to the host where they are created)

27.8 test identifiers

Also known as TGUIDs (or test globally unique identifiers), a test identifier refers to a single unit of test operation or result from the test system. A test identifier may refer to a testcase or an individual test measure.

They consist of a several parts, some of which may be omitted in some circumstances

The parts are:

- testsuite name
- testset name
- testcase name
- measure name

Legal characters for these parts are letters, numbers, and underscores. Only testset names may include a period (“.”), as that is used as the separator between constituent parts of the identifier.

testcase identifiers should be consistent from run-to-run of a test, and should be globally unique.

test identifiers may be in fully-qualified form, or in shortened form - missing some parts. The following rules are used to convert between from shortened forms to fully-qualified forms.

If the testsuite name is missing, then the base name of the test is used.

- e.g. Functional.jpeg has a default testsuite name of “jpeg”

If the testset name is missing, then the testset name is “default”.

A test id may refer to one of two different items:

- a testcase id
- a measure id

A fully qualified test identifier consists of a testsuite name, testset name and a testcase name. Shortened names may be used, in which case default values will be used for some parts, as follows:

If a result id has only 1 part, it is the testcase name. The testset name is considered to be *default*, and the testsuite name is the base name of the test.

That is, for the fuego test Functional.jpeg, a shortened tguid of *test4*, the fully qualified name would be:

- jpeg.default.test4

If a result id has 2 parts, then the first part is the testset name and the second is the testcase name, and the testsuite name is the base name of the test.

27.8.1 measure id

A measure identifier consists of a testsuite id, testset id, testcase id and measure name.

A shortened measure id may not consist of less than 2 parts. If it only has 2 parts, the first part is the testcase id, and the second part is the measure name. In all cases the last part of the name is the measure name, the second-to-last part of the name is the testcase name.

If there are three parts, the first part is the testset name.

27.9 Test variable names

Test variable names are defined in the board file, and by the user via ‘ftc set-var’. Also they are generated from variables declared in spec files. The consist of all upper-case, using only letters and underscores

Some test variable prefixes and suffixes are used in a consistent way.

27.9.1 Dependency check variables

The following is the preferred format for variables used in dependency checking code:

- **PROGRAM_FOO** - require program ‘foo’ on target. The program name is upper-cased, punctuation or spaces are replaced with ‘_’, and the name is prefixed with ‘PROGRAM_’. The value of variable is full path on target where program resides.
 - ex: PROGRAM_BC=/usr/bin/bc
- **HEADER_FOO** - require header file ‘foo’ in SDK. The header filename is stripped of its suffix (I don’t know if that’s a good idea or not), upper-cased, punctuation or spaces are replaced with ‘_’, and the name is prefixed with ‘HEADER_’. The value of variable is the full path in the SDK of the header file:
 - ex: HEADER_FOO=/opt/poky2.1.2/sysroots/x86_64-pokysdk- linux/usr/include/foo.h
- **SDK_LIB_FOO** - require ‘foo’ library in SDK. The library filename is stripped of the ‘lib’ prefix and .so suffix, upper-cased, punctuation and spaces are replaced with ‘_’, and the name is prefixed with ‘SDK_LIB_’. The value of the variable is the full path in the SDK of the library.
 - ex: SDK_LIB_FOO=/opt/poky2.1.2/sysroots/x86_64-pokysdk- linux/usr/lib/libfoo.so
 - Note that in case a static library is required (.a), then the variable name should include that suffix:
 - ex: SDK_LIB_FOO_A=/opt/poky1.2.1/sysroots/x86_64-pokysdk- linux/usr/lib/libfoo.a

- **TARGET_LIB_FOO** - require 'foo' library on target. The library filename is stripped of the 'lib' prefix and .so suffix (not sure this is a good idea, as we potentially lose a library version requirement), upper-cased, punctuation and spaces are replaced with '_', and the name is prefixed with 'TARGET_LIB_'. The value of the variable is the full path of the library on the target board.
 - ex: TARGET_LIB_FOO=/usr/lib/libfoo.so

ARTWORK

This page has artwork (logos, photos and images) for use in Fuego presentations and documents.

28.1 Logos

- fuego-logo.svg (SVG, high-res)



- fuego-logo.png (png, 660x781 w/ transparent background)



- fuego-flame-only.png: (png, 490x617 w/ transparent background)



- fuego-logo-from-Daniel.svg (SVG, high-res)

Please see this image in the Fuegotest wiki at: <http://fuegotest.org/files/fuego-logo-from-Daniel.svg>
(Sphinx cannot include this image in PDF output)

28.2 Banners

- Fuego Jamboree: (png, 567x120)



- firestrip: (jpg, 1710x282)



28.3 images

- vertical fire (public domain, png, 687x2712)



- horizontal fire (public domain, jpg, 615x318)



- rectangle fire (public domain, jpg, 3840x2160)

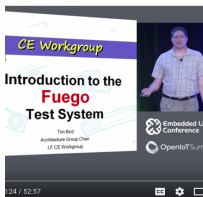


- rectangle fire with Fuego (png, 1431x1056)



28.4 Photos

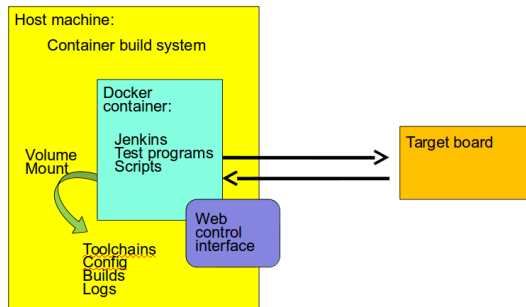
- Tim Bird presenting Introduction to Fuego at ELC 2016 (Youtube video framegrab) (png, 370x370)



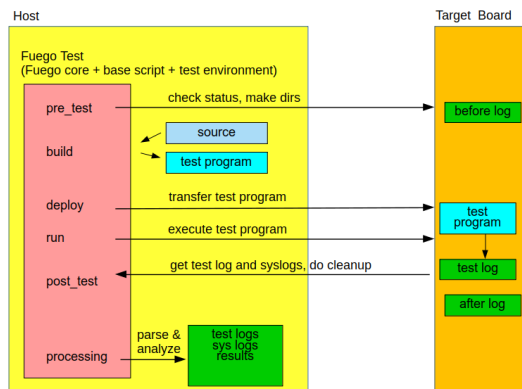
28.5 Diagrams

- Fuego architecture (png, 811x566)

Architecture Diagram



- Fuego test phases (png,) {{BR}}



28.6 Presentation templates

There are templates for making Fuego presentations, on the Fuego wiki at: http://fuegotest.org/wiki/Presentation_Templates

Here is a list of Frequently Asked Questions and Answers about Fuego:

29.1 Languages and formats used

29.1.1 Q. Why does Fuego use shell scripting as the language for tests?

There are other computer languages which have more advanced features (such as data structures, object orientation, rich libraries, concurrency, etc.) than shell scripting. It might seem odd that shell scripting was chosen as the language for implementing the base scripts for the tests in fuego, given the availability of these other languages.

The Fuego architecture is specifically geared toward host/target testing. In particular, tests often perform a variety of operations on the target in addition to the operations that are performed on the host. When the base script for a test runs on the host machine, portions of the test are invoked on the target. It is still true today that the most common execution environment (besides native code) that is available on almost every embedded Linux system is a POSIX-compliant shell. Even devices with very tight memory requirements usually have a busybox ‘ash’ shell available.

In order to keep the base script consistent, Fuego uses shell scripting on both the host and target systems. Shell operations are performed on the target using ‘cmd’, ‘report’ and ‘report_append’ functions provided by Fuego.

Note that Fuego officially use ‘bash’ as the shell on the host, but does not require a particular shell implementation to be available on the target. Therefore, it is important to use only POSIX-compatible shell features for those aspects of the test that run on target.

GLOSSARY

Here is a glossary of terms used in this wiki:

Here is a short table that relates a few Jenkins terms to Fuego terms:

Jenkins term	Fuego term	Description
slave	“none”	this is a long-running jenkins process, that executes jobs. It is usually (?) assigned to a particular node
node	board	item being tested (Fuego defines Jenkins node for each board in the system)
job	test	a collection of information needed to perform a single test
“none”	request	a request to run a particular test on a board
build	run(or ‘test run’)	the results from executing the job or test
“none”	plan	the plan has the list of tests and how to run them (which variation, or ‘spec’ to use)
“none”	spec	the spec indicates a particular variation of a test

30.1 B

base test script

See `test script`.

benchmark

A type of test where one or more numeric metrics indicates the status of the test. See Benchmark parser notes for more information about processing these metrics.

binary package

A tarfile containing the materials that would normally be deployed to the board for execution.

board file

A file that describes (using environment variables) the attributes of a target board. This has the extension `.board` and is kept in the directory `/fuego-ro/boards`.

build

In Jenkins terminology, a “build” is an execution of a test, and the resulting artifacts associated with it. In Fuego, this is also referred to as a test “run”.

30.2 C

`console log`

The full output of execution of a test from Jenkins. See *Log files* for details.

`criteria`

This is an expression which indicates whether how a test result should be interpreted. This is also referred to as a “pass criteria”. Criteria files are stored in `criteria.json` files. See *criteria.json* for details.

30.3 D

`devlog`

The developer log for a test. See *Log files* for details.

`Device`

The name of a target board in the Fuego system.

`device under test`

In test terminology, this refers to the item being tested. In Fuego, this may also be called the “Device”, “board”, “node”, or `target`

`distribution`

This refers to a set of software that is running on a Linux machine. Example “distributions” are Debian, Angstrom or Android. The distribution defines file locations, libraries, utilities and several important facilities and services of the machine (such as the init process, and the logger).

30.4 F

`functional test`

A type of test that returns a single pass/fail result, indicating whether the device under test. It may include lots of sub-tests.

30.5 J

`Jenkins`

An advanced continuous integration system, used as the default front-end for the Fuego test framework. see `Jenkins`

`job`

In Jenkins terminology, a job is a test

30.6 L

log file

Several log files are created during execution of a test. For details about all the different log files, see [Log files](#).

30.7 M

metric

A numeric value measured by a benchmark test as the result of the test. This is compared against a threshold value to determine if the test passed or failed. See Benchmark parser notes

30.8 O

overlay

This is a set of variables and functions stored in a fuegoclass file, which are used to customize test execution for a particular board. See Overlay Generation for details.

ovgen.py

Program to collect “overlay” data from various scripts and data files, and produce the final test script to run. see Overlay Generation.

30.9 P

package

See test package.

parsed log

The test log file after it has been filtered by log_compare. See [Log files](#) for details.

parser.py

A python program, included with each Benchmark test, to scan the test log for benchmark metrics, check each against a reference threshold, and produce a plot.png file for the test. See [parser.py](#) and Benchmark parser notes for more information.

provision

To provision a board is to install the system software on it. Some board control systems re-provision a board for every test. In general, Fuego runs a series of tests with a single system software installation.

30.10 R

reference log

This file (called “reference.log”) defines the regression threshold (and operation) for each metric of a benchmark test. See `reference.log` and Benchmark parser notes

run

See `test run`.

30.11 S

spec variable

A test variable that comes from a spec file. See *Test variables*

stored variable

A test variable that is stored in a read/write file, and can be updated manually or programmatically. See *Test variables*

syslog

The system log for a test. This is the system log collected during execution of a test. See *Log files* for details.

30.12 T

test

This is a collection of scripts, jenkins configuration, source code, and data files used to validate some aspect of the device under test. See Fuego Object Details for more information.

test log

This is the log output from the actual test program on the target. There are multiple logs created during the execution of a test, and some might casually also be called “test logs”. However, in this documentation, the term “test log” should be used only to refer to the test program output. See *Log files* for details.

test package

This is a packaged version of a test, including all the materials needed to execute the test on another host. See Test package system

test phases

Different phases of test execution defined by Fuego: `pre_test`, `build`, `deploy`, `test_run`, `get_testlog`, `test_processing`, `post_test`. For a description of phases see: fuego test phases

test program

A program that runs on the target to execute a test and output the results. This can be a compiled program or a shell script (in which case the build step is empty)

test run

This is a single instance of a test execution, containing logs and other information about the run. This is referred to in Jenkins as a ‘build’.

`test script`

The shell script that interfaces between the Fuego core system and a test program. This is a small script associated with each test. It is called `fuego_test.sh`, and it provides a set of test functions that are executed on the host (in the container) when a test is run.

The script declares a tarfile, and functions to build, deploy and run the test. The test script runs on the host. This is also called the ‘base test script’. For details about the environment that a script runs in or the functions it may call, see [Variables](#), [Core interfaces](#), and Test Script APIs.

`test variable`

This is the name of a variable available to the a test during it’s execution. See [Test variables](#).

`TOOLCHAIN`

Defines the toolchain or SDK for the device. This is used to select a set of environment variables to define and configure the toolchain for building programs for the intended test target.

`tools.sh`

File containing the definition of toolchain variables for the different platforms installed in the container (and supported by the test environment) See [tools.sh](#) for details.

30.13 V

`variable`

See `test variable`

RASPBERRY PI FUEGO SETUP

This is a list of instructions for setting up a Raspberry Pi board for use with Fuego. These instructions will help you set up the ssh server, used by Fuego to communicate with the board, and the test directory on the machine, that Fuego will use to store programs and files during a test.

These instructions and the screen shots are for a Raspberry Pi Model 3 B, running “Raspbian 9 (stretch)”.

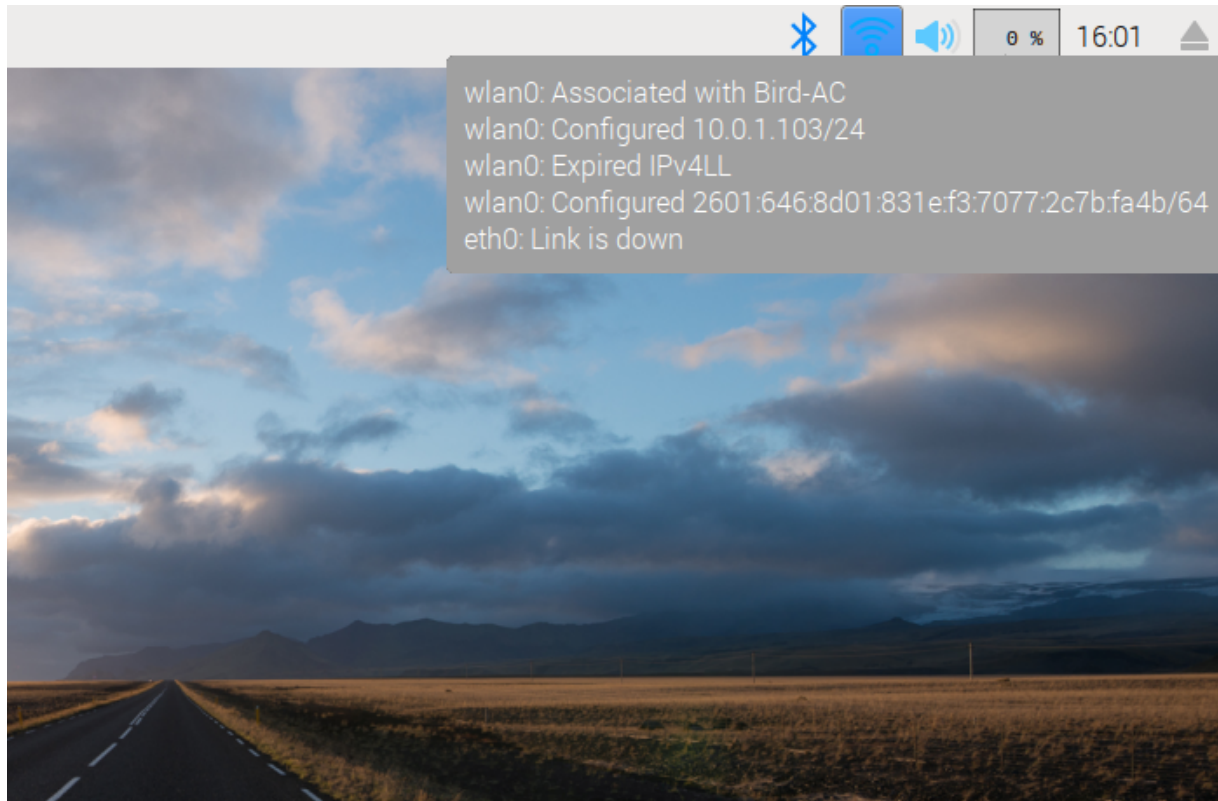
This assumes that the Raspberry Pi is already installed, and that networking is already configured and running.

31.1 Obtain your network address

First, determine what your Pi’s network address is. You can see this by using the command ‘ifconfig’ in a terminal window, and checking for the ‘inet’ address.

Or, move your mouse cursor over the network icon in the desktop panel bar. If you leave the mouse there for a second or two, a box will appear showing information about your current network connection.

This is what the network information box looks like (in the upper right corner of this screen shot):



In this case, my network address is 10.0.1.103. Your address might start with 192.168, which is common for home or local networks.

Note this address for use later.

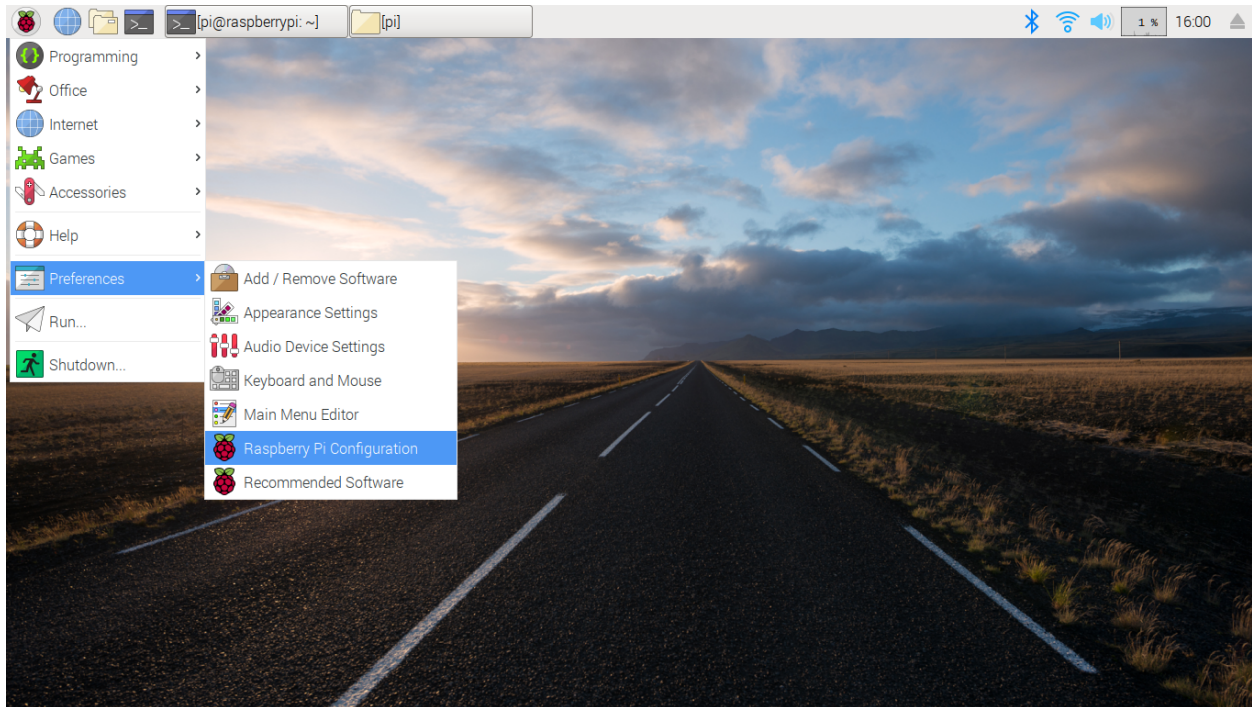
31.2 Configure the SSH server

In order for other machines to access the Pi remotely, you need to enable the ssh server.

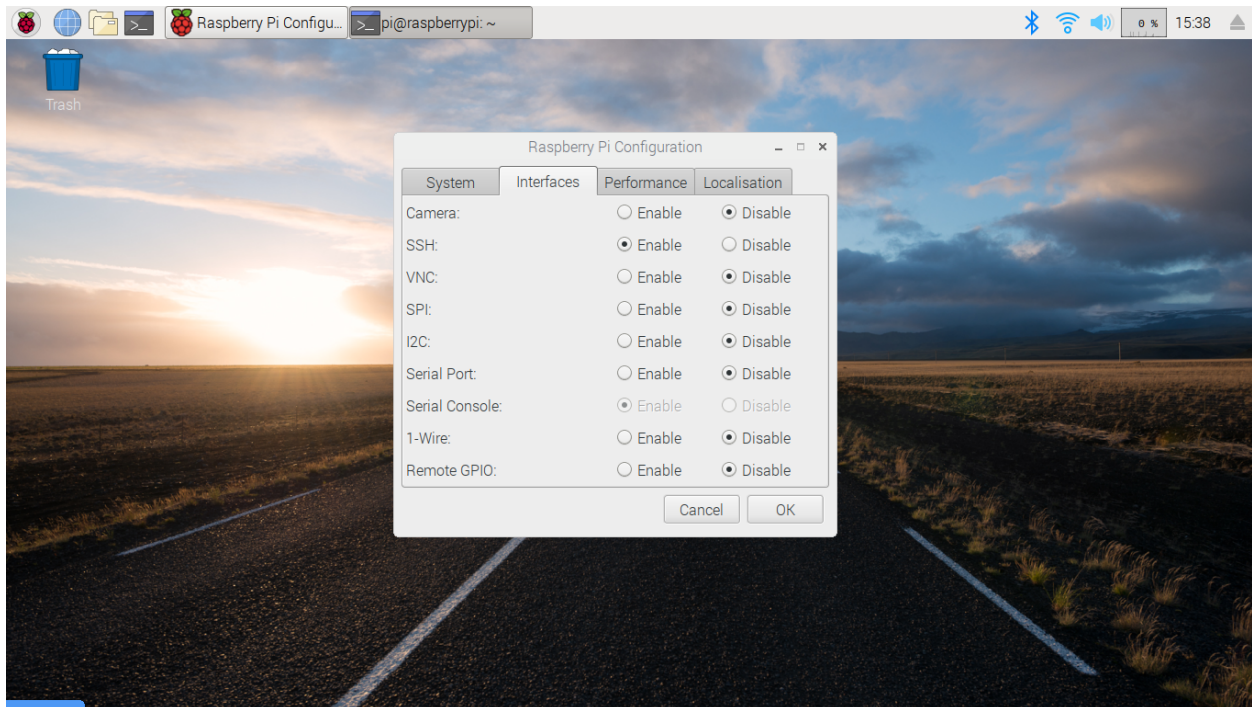
This is done by enabling the SSH interface in the Raspberry Pi Configuration dialog.

To access this dialog, click on the raspberry logo in the upper right corner of the main desktop window. Then click on “Preferences”, then on “Raspberry Pi Configuration”. In the dialog that appears, click on the “Interfaces” tab, and on the list of interfaces click on the “Enable” radio button for the SSH interface.

Here is the menu:



The configuration dialog looks something like this:



31.2.1 Try connecting

Now, close this dialog, and make sure you can access the Pi using SSH from your host machine.

Try the following command, from your host machine:

- `ssh pi@<"your_address">`

You will be asked for the password for the ‘pi’ user account.

If you successfully log in, you will be at a shell prompt.

31.2.2 Configure sshd for root access (if applicable)

If you intend to execute Fuego tests as root, you should configure the SSH server to permit root login.

This is not recommended on machines that are in production, as it is a significant security risk. However, for test machines it may be acceptable to allow root access over ssh.

To do this, on the Raspberry Pi, with root permissions, edit the file `/etc/ssh/sshd_config` and add the following line:

```
PermitRootLogin yes
```

Note: You need to stop and start the ssh server, or reboot the board, in order to have this new configuration take effect.”

31.3 Make a test directory

You can use any directory you like for executing tests from Fuego. However, we recommend using the `/home/fuego` directory. These instructions indicate how to create that directory.

If you are using root as your test user account, then create a directory on the Raspberry Pi for test programs and files.

If logged in as the ‘pi’ account, then switch to root (using something like ‘`sudo su`’), and type, at a shell prompt:

```
$ mkdir /home/fuego
```

If you do not wish to use the root account for testing, then you should create a ‘fuego’ account for testing. To do this, use the ‘`adduser`’ program. You will be prompted for some information.

```
$ adduser fuego
```

Answer the questions, including setting the password for this account. Remember the password you select, and use that in the board file when configuring Fuego to access this board.

This will create the directory `/home/fuego`.

The directory `/home/fuego` is what should be used in the board file when configuring Fuego to operate with this board.

31.4 Add the board file to Fuego

Now that you have set up the Raspberry Pi board, add the board file to Fuego. Assuming your IP address is 10.0.1.17, and you wish to log in as root, you would create a board file called “rpi.board”, and located at

```
<fuego-top-dir>/fuego-ro/boards/rpi.board
```

with the following contents:

```
inherit "base-board"
include "base-params"

IPADDR="10.0.1.17"
LOGIN="root"
BOARD_TESTDIR="/home/fuego"
PASSWORD="put-the-root-password-here"
TOOLCHAIN="debian-armhf"
TRANSPORT="ssh"
ARCHITECTURE="arm"
FUEGO_TARGET_TMP="/home/fuego"
```

Note: Of course, use the correct root password for your board

31.5 Add the toolchain to Fuego

The Raspberry Pi board is an ARM 32-bit platform.

Add the ‘debian-armhf’ toolchain to the Fuego docker container, using the helper script in the fuego-ro/toolchains directory.

Inside the Fuego container, run:

```
$ /fuego-ro/toolchains/install_cross_toolchain.sh armhf
```

31.6 Add a node and jobs for the board

Inside the Fuego container, run:

```
$ ftc add-node -b rpi
```

Add the tests you want to run, as Jenkins jobs. You should always add the “fuego_board_check” test, as a way to automatically determine that status of a board.

Inside the Fuego container, run:

```
$ ftc add-job -b rpi -t Functional.fuego_board_check
```

An easy way to populate Jenkins with a set of tests is to install a batch test.

Install the “smoketest” batch test, as follows:

Inside the Fuego container, run:

```
$ ftc add-jobs -b rpi -t Functional.batch_smoketest
```

31.7 Run a board check

To see if everything is set up correctly, execute the test: `Functional.fuego_board_check`.

In the Jenkins interface, select “`rpi.default.Functional.fuego_board_check`” and select the menu item “Build Now” on the left hand side of the screen.

Wait a few moments for the test to complete. when the test completes, check the log for the test by clicking on the link to the ‘testlog’.

USING THE QEMUARM TARGET

Here are some quick instructions for using the qemu-arm “board” that is preinstalled in fuego.

Fuego does not ship with a qemuarm image in the repository, but assumes that you have built one with the Yocto Project.

If you don’t have one lying around, you will need to build one. Then you should follow the other steps on this page to configure it to run with Fuego.

32.1 Build a qemuarm image

Here are some quick steps for building a qemuarm image using the Yocto Project: (See the [Project Quick Start](#), for more information)

Note that these steps are for Ubuntu.

- Make sure you have required packages for building the software

```
$ sudo apt-get install gawk wget git-core diffstat unzip texinfo \
gcc-multilib build-essential chrpath socat libsdl1.2-dev xterm
```

- Install the qemu software

```
$ sudo apt-get install qemu-user
```

- Download the latest stable release of the Yocto Project

```
$ git clone git://git.yoctoproject.org/poky
```

- Configure Yocto Project for building the qemuarm target

```
$ cd poky
$ source oe-init-build-env build-qemuarm build-qemuarm
$ edit conf/local.conf
```

- Under the comment about “Machine Selection”, uncomment the line

```
MACHINE ?= "qemuarm"
```

- Build a minimal image (this will take a while)

```
$ bitbake core-image-minimal
```

32.2 Running the qemuarm image

You can run the emulator, using the image you just built:

- Run the emulator

```
$ runqemu qemuarm
```

- Find the address and ssh port for the image
 - Inside the image, do `ifconfig eth0`

32.3 Test connectivity

From the host, verify that the networking is running:

```
$ ping 192.168.7.2
$ ssh root@192.168.7.2
```

Of course, substitute the correct IP address in the commands above.

Once you know that things are working, directly connecting from the host to the qemuarm image, make sure the correct values are in the `qemu-arm.board` file. You can edit this file on your host in `fuego-ro/boards/qemu-arm.board`

Here are the values you should set:

```
IPADDR="192.168.7.2"
SSH_PORT=22
LOGIN="root"
PASSWORD=""
```

32.4 Test building software

It is important to be able to build the test software for the image you are using with qemu.

The toolchain used to compile programs for a board is controlled via the `TOOLCHAIN` variable in the board file. Currently the `qemu-arm.board` file specifies `TOOLCHAIN="debian-armhf"`.

You may need to install the `debian-armhf` toolchain, or your own SDK from your Yocto Project build, into the Fuego container in order to build test programs for the qemuarm emulator. See [Adding a toolchain](#) for information about how to do that.

Try building a simple program, like `hello_world`, as a test for the new system, and see what happens. You can try this by executing the test `Functional.hello_world`. You can do that with this command:

```
$ ftc run-test -b qemu-arm -t hello_world
```

VARIABLES

This is an index of all the variables used by Fuego:

FIXTHIS - I don't have all the fuego variables documented here yet

See also *Core interfaces*

33.1 A

- ARCHITECTURE : the processor architecture of the target board
 - Defined in the board file for a target
 - Used by toolchain and build scripts for building the tests
 - * *NOTE: this appears to only be used by iozone.sh*
 - Example value: **arm**
- ARCH : architecture used by the toolchain
 - Example value: **arm**
 - Set by *tools.sh* based on TOOLCHAIN
- AS : name of the assembler
 - Set by *tools.sh* based on TOOLCHAIN
 - Commonly used during the build phase (in the function test_build)

33.2 B

- BAUD : Baud rate to be used with the serialport
 - Defined in the board file for a target
 - Used by serial transport
 - Example value: **115200**
- BOARD_TESTDIR : directory on the target board where test data will be placed
 - Defined in the board file for a target
 - Example value: **/home/fuego**

33.3 C

- CC : name of the C compiler
 - Set by *tools.sh* based on TOOLCHAIN
 - Commonly used during the build phase (in the function test_build)
 - Example value: **arm-linux-gnueabi-gcc**
- CONFIGURE_FLAGS : flags used with the ‘configure’ program
 - Set by *tools.sh* based on TOOLCHAIN
 - Commonly used during the build phase (in the function test_build)
- CROSS_COMPILE : cross-compile prefix used for kernel builds
 - Set by *tools.sh* based on TOOLCHAIN
 - Example value: **arm-linux-gnueabi-**
 - *NOTE: this is often \$PREFIX followed by a single dash*
- CPP : name of the C pre-processor
 - Set by *tools.sh* based on TOOLCHAIN
- CXX : name of the C++ compiler
 - Set by *tools.sh* based on TOOLCHAIN
- CXXCPP : name of the C++ pre-processor
 - Set by *tools.sh* based on TOOLCHAIN

33.4 F

- FUEGO_BUILD_FLAGS : has special flags used to control builds (for some tests)
 - See *FUEGO BUILD FLAGS*
- FUEGO_CORE : directory for Fuego core scripts and tests
 - This is defined in Jenkins and Fuego system-level configurations
 - This will always be /fuego-core inside the Docker container, but will have a different value outside the container.
 - Example value: **/fuego-core**
- FUEGO_DEBUG : controls whether Fuego emits debug information during test execution. This variable is now deprecated in favor of FUEGO_LOGLEVELS
 - See *FUEGO DEBUG*
- FUEGO_LOGLEVELS : controls what level of messages Fuego emits during test execution
 - See *FUEGO LOGLEVELS*
- FUEGO_RO : directory for Fuego read-only data
 - This is defined in the Jenkins and Fuego system-level configurations
 - This will always be /fuego-ro inside the Docker container, but will have a different value outside the container.

- Example value: **/fuego-ro**
- FUEGO_RW : directory for Fuego read-write data
 - This is defined in Jenkins and Fuego system-level configurations
 - This will always be **/fuego-rw** inside the Docker container, but will have a different value outside the container.
 - Example value: **/fuego-rw**
- FUEGO_TARGET_TMP : directory on target to use for syslogs
 - This is defined in the board file for a target board
 - This should specify a directory in the board filesystem that is persistent across reboots. This is to override the default temp directory (of **/tmp**), if that directory is erased on a board reboot.
- FUEGO_TEST_PHASES : specifies a list of phases to perform during test execution
 - This is usually set by **ftc run-test** based on the **‘-p’** option.
 - This is a space-separated list of strings, from the following possible individual phase strings: **pre_test**, **pre_check**, **build**, **deploy**, **snapshot**, **run**, **post_test**, **processing**, **makepkg**
 - Example value: **pre_test pre_check build deploy snapshot run post_test processing**

33.5 G

- GEN_TESTRES_FILE : set to the value of TEST_RES, when a BATCH_TESTPLAN is in effect

33.6 I

- IO_TIME_SERIAL : Time required for echoing the whole command and response
 - Defined in the board file
 - Used by the transport functions
 - Example value: **0.1**
- IPADDR : IP address of the target board
 - Defined in the board file
 - Used by the transport functions
 - Example value: **10.0.1.74**

33.7 L

- LD : name of the linker
 - Set by *tools.sh* based on TOOLCHAIN
 - Example value: **arm-linux-gnueabi-hf-ld**
- LOGIN : login account name for the target
 - Defined in the board file for the target
 - Used by the transport functions
 - The account on the target should have sufficient rights to run a variety of tests and perform a variety of operations on the target
 - Example value: **root**

33.8 M

- MAX_BOOT_RETRIES : Number of times to retry connecting to target during a target_reboot operation.
 - Defined in the board file
 - Example value: **20**
- MMC_DEV : device filename for an MMC device on the target
 - Defined in the board file
 - Used by filesystem test specs
 - Example value: **/dev/mmcblk0p2**
- MMC_MP : mount point for a filesystem on an MMC device on the target
 - Defined in the board file
 - Used by filesystem test specs
 - Example value: **/mnt/mmc**
- MOUNT_BLOCKDEV : device filename for a block device on the target
 - Defined in a filesystem test spec
 - * e.g. in (bonnie, fio, ffsb, iohzone, synctest, aiostress, dbench, tiobench).spec
 - Usually references either MMC_DEV, SATA_DEV or USB_DEV, depending on what the test spec indicates to test
 - Example value: **/dev/sda1**
- MOUNT_POINT : mount point for a filesystem to be tested on the target
 - Defined in a filesystem test spec
 - * e.g. in (bonnie, fio, ffsb, iohzone, synctest, aiostress, dbench, tiobench).spec
 - Usually references either MMC_MP, SATA_MP, or USB_MP, depending on what the test spec indicates to test
 - Example value: **/mnt/sata**

33.9 N

- **NO_BOARD_TESTLOG** : indicates that the test does not produce testlog output on the board
 - Set ‘**NO_BOARD_TESTLOG=1**’ in your **fuego_test.sh** file to tell Fuego that this test does not produce test output on the board. This will suppress a message about a missing log file, after test execution.
- **NODE_NAME** : the name of the board
 - This is set by Jenkins, and is the first part of the Fuego job name

33.10 O

- **OF_ROOT** : root of overlay system
 - Example value: **/home/jenkins/overlays/**

33.11 P

- **PASSWORD** : password used with to login to the target board
 - Defined in the board file for a target
 - Used by the transport functions
 - It can be the empty string: “”
 - Example value: **mypass**
- **PLATFORM** : name of the target “platform” This is used to identify the toolchain used for building tests. This has been deprecated. Please use **TOOLCHAIN** instead.
- **PREFIX** : toolchain prefix
 - Set by *tools.sh* based on **TOOLCHAIN**
 - Example value: **arm-linux-gnueabi**
 - *NOTE: see also CROSS_COMPILE*

33.12 R

- **REP_DIR** : directory where reports are stored
 - Example value: **/home/jenkins/logs/logruns/**
- **REP_GEN** : report generator program
 - Example value: **/home/jenkins/scripts/loggen/gentxml.py**
- **REP_LOGGEN** : program used to generate report logs?
 - Example value: **/home/jenkins/scripts/loggen/loggen.py**

33.13 S

- **SATA_DEV** : device node filename for a SATA device on the target
 - Defined in the board file
 - Used by filesystem tests
 - Example value: **/dev/sda1**
- **SATA_MP** : mount point for a filesystem on a SATA device on the target
 - Used by filesystem tests
 - Example value: **/mnt/sata**
- **SRV_IP** : IP address of server machine (host machine where fuego runs)
 - Defined in base-board.fuegoclass
 - * Obtained dynamically using the **ip** command
 - Can be defined in a board file for a target, using an **override** command
 - Used by networking tests (NetPIPE, iperf, netperf)
 - Example value: **10.0.1.1**
- **SSH_PORT** : port to use for ssh connections on the target
 - Defined in the board file for the target
 - The default port for sshd is 22
 - Example value: **22**
- **SERIAL** : port to use for serial connections on the target
 - Defined in the board file for the target
 - The device file name as detected in Docker container
 - Example value: **ttyACM0**

33.14 T

- **TESTLOG** : full path to log for a particular test
 - Example value: **/home/jenkins/logs/Functional.bzip2/ testlogs/bbb.2016-06-24_18-12-53.2.log**
- **TEST_RES** : full path to JSON results file for a test
 - Example value: **/home/jenkins/logs/Functional.bzip2/testlogs/ bbb.2016-06-24_18-12-53.2.res.json**
- **TESTDIR** : name of the directory for a particular test
 - This is just the directory name, not the full path (see \$TEST_HOME)
 - This is also used as the reference parse log prefix
 - Example value: **Functional.bzip2**
- **TEST_HOME** : full path to the root of the test directory
 - Example value: **/fuego-core/tests/Functional.bzip2**
- **TOOLCHAIN** : name of the toolchain used to build test programs for a board.

- Defined in the board file
- Used in `tools.sh`
- Example value: **debian-armhf**
- *NOTE: this replaced 'PLATFORM', used in earlier versions of Fuego*
- **TRANSPORT** : type of connection between the host system and the target system
 - Defined in the board file for the target
 - possible values: **ssh**, **serial**, **ttc**, **ssh2serial**, **local**
 - * Others anticipated are: **adb**, **lava**
 - Used by the transport functions
 - Example value: **ssh**
- **TTC_TARGET** : target name used with `ttc` command
 - Defined in the board file for the target
 - Used by the transport functions, for the `ttc` transport only
 - Example value: **beaglebone**

33.15 U

- **USB_DEV** : device filename for an block device provided by a USB device on the target
 - Defined in the board file
 - Used by filesystem test specs
 - Example value: **/dev/sdb1**
- **USB_MP** : mount point for a filesystem on an USB device on the target
 - Defined in the board file
 - Used by filesystem test specs
 - Example value: **/mnt/usb**

33.16 UNDOCUMENTED (YET)

- **TRIPLET**
- **LTP_OPEN_POSIX_SUBTEST_COUNT_POS**
 - Defined in board file for a target
- **LTP_OPEN_POSIX_SUBTEST_COUNT_NEG**
 - Defined in board file for a target
- **EXPAT_SUBTEST_COUNT_POS**
 - Defined in board file for a target
- **EXPAT_SUBTEST_COUNT_NEG**
 - Defined in board file for a target

- OF_ROOT
- OF_CLASSDIR
- OF_DEFAULT_SPECDIR
- OF_OVFILES
- OF_CLASSDIR_ARGS
- OF_TESTPLAN_ARGS
- OF_SPECDIR_ARGS
- OF_OUTPUT_FILE
- OF_OUTPUT_FILE_ARGS
- OF_DISTRIB_FILE
- OF_OVGEN
- OF_BOARD_FILE
- BATCH_TESTPLAN
- OF_TESTPLAN
- OF_TESTPLAN_ARGS
- OF_OVFILES_ARGS

DYNAMIC VARIABLES

“Dynamic variables” in Fuego are variables that can be passed to a test on the command line, and used to customize the operation of a test, for a particular test run.

In general testing nomenclature this is referred to as test parameterization.

The purpose of dynamic variables is to support “variant” testing, where a script can loop over a test multiple times, changing the variable to different values.

In Fuego, during test execution dynamic variable names are expanded to full variables names that are prefixed with the name of the test. A dynamic variable overrides a spec variable of the same name.

Here is an example of using dynamic variables:

```
$ ftc run-test -b beaglebone -t Benchmark.Dhrystone --dynamic-vars "LOOPS=100000000"
```

This would override the default value for `BENCHMARK_DHRYSTONE_LOOPS`, setting it to 100000000 (100 million) for this run. Normally, the default spec for `Benchmark.Dhrystones` specifies a value of 10000000 (10 million) for `LOOPS`.

This feature is intended to be useful for (among other things) doing ‘git bisect’s of a bug, passing a different git commit id for each iteration of the test.

See `Test_variables` for more information.

34.1 Notes

Note that dynamic vars are added to the runtime `spec.json` file, which is saved in the log directory for the run being executed.

This `spec.json` file is copied from the one specified for the run (usually from the test’s home directory).

If dynamic variables have been defined for a test, then they are listed by name in the run-specific `spec.json` file, as the value of the variable “`dyn_vars`”. The reason for this is to allow someone who reviews the test results later to easily see whether a particular test variable had a value that derived from the spec, or from a dynamic variable. This is important for proper results interpretation.

TROUBLESHOOTING GUIDE

This page describes problems encountered using Fuego, and their solutions.

Note: for Editors: please put each issue in it's own page section.

35.1 Installation

35.1.1 Problem with default Jenkins port

Fuego has Jenkins default to using port 8090 on the host system. However, if you have something else already running on port 8090, you may wish to change this.

You can change the Jenkins port during installation of Fuego, using an argument to the `install.sh` script. For example, to install Fuego with Jenkins configured to use port 9999, use the following command during installation:

```
$ ./install.sh fuego 9999
```

To change the Jenkins port for an already-built Fuego container, start the container, and inside the container edit the file:

- `/etc/default/jenkins`

Change the line that says: `HTTP_PORT=8090`

Change to port to whatever your like.

Also, check the line that defines `JENKINS_ARGS`. Mine looked like this:

```
JENKINS_ARGS="--webroot=/var/cache/jenkins/war --httpPort=8090 --prefix=/fuego"
```

Change this line to read as follows:

```
JENKINS_ARGS="--webroot=/var/cache/jenkins/war --httpPort=$HTTP_PORT --prefix=/fuego"
```

Then restart Jenkins:

- `$ service jenkins restart`

35.1.2 Problem creating docker file

Make sure you are running on a 64-bit version of the Linux kernel on your host machine.

35.1.3 Problem starting Jenkins after initial container creation

Doug Crawford reported a problem starting Jenkins in the container after his initial build.

```
$ sudo ./docker-create-container.sh
Created JTA container 6a420f901af7847f2afa3100d3fb3852b71bc65f92aecdd13a9aefe0823d42b77
$ sudo ./docker-start-container.sh Starting JTA container
6a420f901af7847f2afa3100d3fb3852b71bc65f92aecdd13a9aefe0823d42b77
[....] Starting Jenkins Continuous Integration Server: jenkins: System error failed!
[ ok ] Starting OpenBSD Secure Shell server: sshd.
[ ok ] Starting network benchmark server.
```

The error string is jenkins: System error

Takuo Kogushi provides the following response:

I had the same issue. I did some search in the net and found it is not a problem of fuego itself. As far as I know there are two workarounds;

- 1) Rebuild and install libpam with `--disable-audit` option (in the container) or
- 2) Modify `docker-create-container.sh` to add `--pid="host"` option to docker create command

Here is a patch provided by Koguchi-san:

```
diff --git a/fuego-host-scripts/docker-create-container.sh b/fuego-host-scripts/docker-
↪create-container.sh
index 2ea7961..24663d6 100755
--- a/fuego-host-scripts/docker-create-container.sh
+++ b/fuego-host-scripts/docker-create-container.sh
@@ -7,7 +7,7 @@ while [ -h "$SOURCE" ]; do # resolve $SOURCE until the file is no longer_
↪a symli done DIR="$( cd -P "$( dirname "$SOURCE" )" && pwd )"

-CONTAINER_ID=`sudo docker create -it -v $DIR/../userdata:/userdata --net="host" fuego`
+CONTAINER_ID=`sudo docker create -it -v $DIR/../userdata:/userdata --pid="host" --net=
↪"host" fuego`
CONTAINER_ID_FILE="$DIR/../last_fuego_container.id"
echo "Created Fuego container $CONTAINER_ID"
echo $CONTAINER_ID > $DIR/../last_fuego_container.id
```

Actually I have not tried the first one and do not know if there is any side effects for the second.

This may be related to this docker bug: <https://github.com/docker/docker/issues/5899>

35.2 General

35.2.1 Timeout executing ssh commands

In some cases, the ssh command used by Fuego takes a very long time to connect. There is a timeout for the ssh commands, specified as 15 seconds in the cogent repository and 30 seconds in the fuegotest repository.

The timeout for ssh commands is specified in the file

- `/fuego-core/scripts/overlays/base/base-params.fuegoclass`

You can change `ConnectTimeout` to something longer by editing the file.

FIXTHIS - make `ConnectTimeout` for ssh connections a board-level test variable

35.2.2 ssh commands taking a long time

Sometimes, even if the command does not time, the SSH operations on the target take a very long time for each operation.

The symptom is that when you are watching the console output for a test, the test stops at the point of each SSH connection to the target.

One cause of long ssh connection times can be that the target ssh server (sshd) is configured to do DNS lookups on each inbound connection.

To turn this off, on the target, edit the file:

- `/etc/ssh/sshd_config`

and add the line:

`UseDNS no`

This line can be added anywhere in the file, but I recommend adding it right after the `UsePrivilegeSeparation` line (if that's there).

35.3 Handling different Fuego Error messages

Here are some Fuego error messages, their meaning, and how to fix them.

35.3.1 “Fuego error: Could not read ‘before’ syslog from board”

This message is found in `syslog.before.txt` in the log directory for a test run. This indicates that Fuego could not find the ‘before’ syslog during test execution. This can happen if a board reboots during a test, and the ‘before’ syslog was stored in a temp directory on the board that is cleared on board reboot. This is not a fatal error.

To avoid this problem, specify a different, non-ephemeral, tmp directory in the board file for the board, using the variable `FUEGO_TARGET_TMP`

Here is an example:

`FUEGO_TARGET_TMP="/home/fuego/tmp"`

OSS TEST VISION

This page describes aspects of vision for the Fuego project, along with some ideas for implementing specific ideas related to this vision.

36.1 Overview of concepts

1. Decentralized testing
2. Automated selection of tests based on test or platform attributes
3. Standardized definition of test attributes and dependencies
4. Way to connect developers with relevant test hardware
5. Test Store (a public repository of available tests)
6. Way to share test-related information (useful parameters, results interpretation)
7. Standards for test packaging
8. Provide incentives for test activities

36.1.1 Letter to ksummit discuss

Here's an e-mail Tim sent to the ksummit-discuss list in October, 2016

I have some ideas on Open Source testing that I'd like to throw out there for discussion. Some of these I have been stewing on for a while, while some came to mind after talking to people at recent conference events.

Sorry - this is going to be long...

First, it would be nice to increase the amount of testing we do, by having more test automation. (ok, that's a no-brainer). Recently there has been a trend towards more centralized testing facilities, like the zero-day stuff or board farms used by *KernelCI*. That makes sense, as this requires specialized hardware, setup, or skills to operate certain kinds of test environments. As one example, an automated test of kernel boot requires automated control of power to a board or platform, which is not very common among kernel developers.

(continues on next page)

(continued from previous page)

A centralized test facility has the expertise and hardware to add new test nodes relatively cheaply. They can do this more quickly and much less expensively than the first such node by an individual new to testing.

However, I think to make great strides in test quantity and coverage, it's important to focus on ease of use for individual test nodes. My vision would be to have tens of thousands of individual test nodes running automated tests on thousands of different hardware platforms and configurations and workloads.

The kernel selftest project is a step in the right direction for this, because it allows any kernel developer to easily (in theory) run automated unit tests for the kernel. However, this is still a manual process. I'd like to see improved standards and infrastructure for automating tests.

It turns out there are lots of manual steps in the testing and bug-fixing process with the kernel (and other Linux-related software). It would be nice if a new system allowed us to capture manual steps, and over time convert them to automation.

Here are some problems with the manual process that I think need addressing:

1) How does an individual know what tests are valid for their platform? Currently, this is a manual decision. In a world with thousands or tens of thousands of tests, this will be very difficult.

We need to have automated mechanisms to indicate which tests are relevant for a platform. Test definitions should include a description of the hardware they need, or the test setup they need. For example, it would be nice to have tests indicate that they need to be run on a node with USB gadget support, or on a node with the gadget hardware from a particular vendor (e.g. a particular SOC), or with a particular hardware phy (e.g. Synopsis). As another example, if a test requires that the hardware physically reboot, then that should be indicated in the test. If a test requires that a particular button be pressed (and that the button be available to be pressed), it should be listed. Or if the test requires that an external node be available to participate in the test (such as a wifi endpoint, CANbus endpoint, or i2C device) be present, that should be indicated. There should be a way for the test nodes which provide those hardware capabilities, setups, or external resources to identify themselves.

Standards should be developed for how a test node and a test can

(continues on next page)

(continued from previous page)

express these capabilities and requirements. Also, standards need to be developed so that a test can control those external resources to participate in tests. Right now each test framework handles this in its own way (if it provides support for it at all).

I heard of a neat setup at one company where the video output from a system was captured by another video system, and the results analyzed automatically. This type of test setup currently requires an enormous investment of expertise, and possibly specialized hardware. Once such a setup is performed in a few locations, it makes much more sense to direct tests that need such facilities to those locations, than it does to try to spread the expertise to lots of different individuals (although that certainly has value also).

For a first pass, I think the kernel CONFIG variables needed by a test should be indicated, and they could be compared with the config for the device under test. This would be a start on the expression of the dependencies between a test and the features of the test node.

2) How do you connect people who are interested in a particular test with a node that can perform that test?

My proposal here is simple - for every subsystem of the kernel, put a list of test nodes in the MAINTAINERS file, to indicate nodes that are available to test that subsystem. Tests can be scheduled to run on those nodes, either whenever new patches are received for that sub-system, or when a bug is encountered and developers for that subsystem want to investigate it by writing a new test. Tests or data collection instructions that are now provided manually would be converted to formal test definitions, and added to a growing body of tests. This should help people re-use test operations that are common. Capturing test operations that are done manually into a script would need to be very easy (possibly itself automated), and it would need to be easy to publish the new test for others to use.

Basically, in the future, it would be nice if when a person reported a bug, instead of the maintainer manually walking someone through the steps to identify the bug and track down the problem, they could point the user at an existing test that the user could easily run.

I imagine a kind of "test app store", where a tester can select from thousands of tests according to their interest. Also, people could rate the tests, and maintainers could point people to tests that are helpful to solve specific problems.

(continues on next page)

(continued from previous page)

3) How does an individual know how to execute a test and how to interpret the results?

For many features or sub-systems, there are existing tools (e.g bonnie for filesystem tests, netperf for networking tests, or cyclicttest for realtime), but these tools have a variety of options for testing different aspects of a problem or for dealing with different configurations or setups. Online you can find tutorials for running each of these, and for helping people interpret the results. A new test system should take care of running these tools with the proper command line arguments for different test aspects, and for different test targets ('device-under-test's).

For example, when someone figures out a set of useful arguments to cyclicttest for testing realtime on a beaglebone board, they should be able to easily capture those arguments to allow another developer using the same board to easily re-use those test parameters, and interpret the cyclicttest results, in an automated fashion. Basically we want to automate the process of finding out "what options do I use for this test on this board, and what the heck number am I supposed to look at in this output, and what should its value be?".

Another issue is with interpretation of test results from large test suites. One notorious example of this is LTP. It produces thousands of results, and almost always produces failures or results that can be safely ignored on a particular board or in a particular environment. It requires a large amount of manual evaluation and expertise to determine which items to pay attention to from LTP. It would be nice to be able to capture this evaluation, and share it with others with either the same board, or the same test environment, to allow them to avoid duplicating this work.

Of course, this should not be used to gloss over bugs in LTP or bugs that LTP is reporting correctly and actually need to be paid attention to.

4) How should this test collateral be expressed, and how should it be collected, stored, shared and re-used?

There are a multitude of test frameworks available. I am proposing that as a community we develop standards for test packaging which include this type of information (test dependencies, test parameters, results interpretation). I don't know all the details yet. For this reason I am coming to the community see how others are solving these problems and to get ideas for how to solve them in a way that would

(continues on next page)

(continued from previous page)

be useful for multiple frameworks. I'm personally working on the Fuego test framework - see <http://fuegotest.org/wiki>, but I'd like to create something that could be used with any test framework.

5) How to trust test collateral from other sources (tests, interpretation)

One issue which arises with this type of sharing (or with any type of sharing) is how to trust the materials involved. If a user puts up a node with their own hardware, and trusts the test framework to automatically download and execute a never-before-seen test, this creates a security and trust issue. I believe this will require the same types of authentication and trust mechanisms (e.g. signing, validation and trust relationships) that we use to manage code in the kernel.

I think this is more important than it sounds. I think the real value of this system will come when tens of thousands of nodes are running tests where the system owners can largely ignore the operation of the system, and instead the test scheduling and priorities can be driven by the needs of developers and maintainers who the test node owners have never interacted with.

Finally,

6) What is the motivation for someone to run a test on their hardware?

Well, there's an obvious benefit to executing a test if you are personally interested in the result. However, I think the benefit of running an enormous test system needs to be de-coupled from that immediate direct benefit. I think we should look at this the same way we look at other crowd-sourced initiatives, like Wikipedia. While there is some small benefit for someone producing an individual page edit, we need to move beyond that to the benefit to the community of the cumulative effort.

I think that if we want tens of thousands of people to run tests, then we need to increase the cost/benefit ratio for the system. First, you need to reduce the cost so that it is very cheap, in all of [time|money|expertise| ongoing attention], to set up and maintain a test node. Second, there needs to be a real benefit that people can measure from the cumulative effect of participating in the system. I think it would be valuable to report bugs found and fixed by the system as a whole, and possibly to attribute positive results to the output provided by individual nodes. (Maybe you could 'game-ify' the operation of test nodes.)

(continues on next page)

(continued from previous page)

Well, if you are still reading by now, I appreciate it. I have more ideas, including more details for how such a system might work, and what types of things it could accomplish. But I'll save that for smaller groups who might be more directly interested in this topic.

To get started, I will begin working on a prototype of a test packaging system that includes some of the ideas mentioned here: inclusion of test collateral, and package validation. I would also like to schedule a "test summit" of some kind (maybe associated with ELC or Linaro Connect, or some other event), to discuss standards in the area I propose.

I welcome any response to these ideas. I plan to discuss them at the upcoming test framework mini-jamboree in Tokyo next week, and at Plumbers (particularly during the 'testing and fuzzing' session) the week following. But feel free to respond to this e-mail as well.

Thanks.

-- Tim Bird

36.2 Ideas related to the vision

36.2.1 Capturing tests easily

- It should be easy to capture a command line sequence, and test the results
- It might be nice to do an automated capture of command output, and format the output into a `clitest` file that can be used as a here document inside a Fuego test script.

36.3 Test definition or attributes

- Do test definitions need to be board-specific
- Elements of test definition:
 - Test dependencies:
 - * Kernel config values needed
 - * Kernel features needed:
 - proc filesystem
 - sys filesystem
 - trace filesystem
 - * Test hardware needed
 - * Test node setup features

- ability to reboot the board
 - ability to soft-reset the board
 - ability to install a new kernel
- * Presence of certain programs on target
 - bc
 - top, ps, /bin/sh, bash?
- Fuego already has:
 - CAPABILITIES?
 - pn and reference logs
 - positive and negative result counts (specific to board)
 - test specs indicate parameters for the test
 - test plans indicate different profiles (method to match test to test environment - e.g. filesystem test with type of filesystem hardware)

36.4 Test app store

It would be nice to have an “Test Store”, similar to an “App Store”, where tests can be made publicly available, and browsed and installed by test developers, based on their needs.

Here are some of the items needed for this project:

- Need a repository where tests can be downloaded
 - similar to a Jenkins plugin repository
 - or similar to a Debian package feed
- Need a client for browsing tests, installing tests, updating tests
- It might be possible to store tests in github, and just refer to different tests in different git repositories?
- It would be nice to have test ratings, including user feedback on tests
- It would be nice to have test metrics (e.g. how many bugs has the test found)

36.5 Authenticating tests

It is important, if you have a public repository of tests, that you introduce an element of trust and authentication to the repository to avoid malicious actions. You may want to have an authority review the test and possibly sign it. An open question would be who would be the trusted authority (Fuego maintainers? This would turn into a bottleneck).

36.6 Test system metrics

It is useful for a test system to provide information about the number of bugs that a test system finds and that get fixed in upstream software. Also, a test system will find bugs in test programs and in itself. These should be noted as well.

TEST SPECS AND PLANS

Note: This page describes (among other things) the standalone test plan feature, which is in process of being converted to a new system. The information is still accurate as of November 2020 (and Fuego version 1.5). However, standalone testplans are now deprecated. The testplan system is being refactored and testplan json data is being integrated into a new batch test system introduced in Fuego version 1.5. See [Using Batch Tests](#) for more information.

37.1 Introduction

Fuego provides a mechanism to control test execution using something called “test specs” and “test plans”. Together, these allow for customizing the invocation and execution of tests in the system. A test spec lists variables, with their possible values for different test scenarios, and the test plan lists a set of tests, and the set of variable values (the spec) to use for each one.

There are numerous cases where you might want to run the same test program, but in a different configuration (i.e. with different settings), in order to measure or test some different aspect of the system. This is often referred to as test “parameterization”. One of the simplest different types of test settings you might choose is whether to run a quick test or a thorough (long) test. Selecting between quick and long is a high-level concept, and corresponds to the concept of a test plan. The test plan selects different arguments, for the tests that this makes sense for.

For example, the arguments to run a long filesystem stress test, are different than the arguments to run a long network benchmark test. For each of these individual tests, the arguments will be different for different plans.

Another broad category of test difference is what kind of hardware device or media you are running a filesystem test on. For example, you might want to run a filesystem test on a USB-based device, but the results will likely not be comparable with the results for an MMC-based device. This is due to differences in how the devices operate at a hardware layer and how they are accessed by the system. Therefore, depending on what you are trying to measure, you may wish to measure only one or another type of hardware.

The different settings for these different plans (the test variables or test parameters) are stored in the test spec file. Each test in the system has a test spec file, which lists different specifications (or “specs”) that can be incorporated into a plan. The specs list a set of variables for each spec. When a testplan references a particular spec, the variable values for that spec are set by the Fuego overlay generator during the test execution.

In general, test plan files are global and have the names of categories of tests.

Note: Note that a spec mentioned in a test plan may not be available every test. In fact the only spec that is guaranteed to be available in every test is the ‘default’ test spec. It is important for the user to recognize which test specs and plan arguments may be suitably used with which tests.

37.2 Test plans

The Fuego “test plan” feature is provided as an aid to organizing testing activities.

There are only a few “real” testplans provided in Fuego (as of early 2019). There is a “default” testplan, which includes a smattering of different tests, and some test plans that allow for selecting between different kinds of hardware devices that provide file systems. Fuego includes a number of different file system tests, and these plans allow customizing each test to run with filesystems on either USB, SATA, or MMC devices.

These test plans allow this selection:

- `testplan_usbstor`
- `testplan_sata`
- `testplan_mmc`

These plans select test specs named: ‘usb’, ‘sata’, and ‘mmc’ respectively.

Fuego also includes some test-specific test plans (for the `Functional.bc` and `Functional.hello_world` tests), but these are there more as examples to show how the test plan and spec system works, than for any real utility.

A test plan is specified by a file in JSON format, that indicates the test plan name, and a list of tests. For each test, it also lists the specs which should be used for that test, when run with this plan. The test plan file should have a descriptive name starting with ‘**testplan_**’ and ending in the suffix ‘.json’, and the file must be placed in the `overlays/testplans` directory.

37.2.1 Example hello world testplan

The test program from the `hello_world` test allows for selecting whether the test succeeds, always fails, or fails randomly. It does this using a command line argument.

The Fuego system includes test plans that select these different behaviors. These test plan files are named:

- `testplan_default.json`
- `testplan_hello_world_fail.json`
- `testplan_hello_world_random.json`

Here is `testplan_hello_world_random.json`

```
{
  "testPlanName": "testplan_hello_world_random",
  "tests": [
    {
      "testName": "Functional.hello_world",
      "spec": "hello-random"
    }
  ]
}
```

37.2.2 Testplan Reference

A testplan can include several different fields, at the “top” level of the file, and associated with an individual test. These are described on the page: [Testplan Reference](#)

37.3 Test Specs

Fuego’s “test spec” system is a mechanism for running Fuego tests in a “parameterized” fashion. That is, you can run the same underlying test program, but with different values for variables that are passed to the test (the test “parameters”, in testing nomenclature). Each ‘spec’ that is defined for a test may also be referred to as a test ‘variant’ - that is, a variation on the basic operation of the test.

Each test in Fuego should have a ‘test spec’ file, which lists different specifications or variants for that test. For each ‘spec’ (or variant), the configuration declares the variables that are recognized by that test, and their values. Every test is required to define a “default” test spec, which is the default set of test variables used when running the test. Note that a test spec is not required to define any test variables, and this is the case for many ‘default’ test specs for tests which have no variants.

The set of variables, and what they contain is highly test-specific. In some cases, a test variable is used to configure different command line options for the test program. In other cases, the variable may be used by `fuego_test.sh` to change how test preparation is done, or to select different hardware devices or file systems for the test to operate on.

The test spec file is in JSON format, and has the name “spec.json”.

The test spec file is placed in the test’s home directory, which is based on the test’s name: `/fuego-core/tests/$TESTNAME/spec.json`

37.3.1 Example

The `Functional.hello_world` test has a test spec that provides options for executing the test normally (the ‘default’ spec), for always failing (the ‘hello-fail’ spec), or for succeeding or failing randomly (the ‘hello-random’ spec)

This test spec file for the ‘hello_world’ test is `fuego-core/tests/Functional.hello_world/spec.json`

Here is the complete spec for this test:

```
{
  "testName": "Functional.hello_world",
  "specs": {
    "hello-fail": {
      "ARG": "-f"
    },
    "hello-random": {
      "ARG": "-r"
    },
    "default": {
      "ARG": ""
    }
  }
}
```

During test execution, the variable `$FUNCTIONAL_HELLO_WORLD_ARG` will be set to one of the three values shown (nothing, ‘-f’ or ‘-r’), depending on which is spec used when the test is run.

In Fuego, the spec to use with a test can be specified multiple different ways:

- as part of the Jenkins job definition
- on the `ftc run-test` command line
- as part of a testplan definition

37.4 Variable use during test execution

Variables from the test spec are expanded by the overlay generator during test execution. The variables declared in the test spec files may reference other variables from the environment, such as from the board file, relating to the toolchain, or from the fuego system itself.

The name of the variable is appended to the end of the test name to form the environment variable that is used by the test. The environment variable name is converted to all uppercase. This environment variable can be used in the `fuego_test.sh` as an argument to the test program, or in any other way desired.

37.4.1 Example of variable use

In this hello-world example, the program invocation (by `fuego_test.sh`) uses the variable `$FUNCTIONAL_HELLO_WORLD_ARG`. Below is an excerpt from `/fuego-core/tests/Functional.hello_world/fuego_test.sh`.

```
function test_run {  
    report "cd $BOARD_TESTDIR/fuego.$TESTDIR; ./hello $FUNCTIONAL_HELLO_WORLD_ARG"  
}
```

Note that in the default spec for `hello_world`, the variable (`'ARG'` in the test spec) is left empty. This means that during execution of this test with `testplan_default`, the program `'hello'` is called with no arguments, which will cause it to perform its default operation. The default operation for the `'hello'` program is to write `"Hello World"` and a test result of `"SUCCESS"`, and then exit successfully.

37.5 Specifying failure cases

A test spec file can also specify one or more failure cases. These represent string patterns that Fuego scans for in the test log, to detect error conditions indicating that the test failed. The syntax for this is described next.

37.5.1 Example of fail case

The following example of a test spec (from the `Functional.bc` test), shows how to declare an array of failure tests for this test.

There should be an variable named `"fail_case"` declared in test test spec JSON file, and it should consist of an array of objects, each one specifying a `'fail_regexp'` and a `'fail_message'`, with an optional variable (`use_syslog`) indicating to search for the item in the system log instead of the test log.

The regular expression is used with `grep` to scan lines in the test log. If a match is found, then the associated message is printed, and the test is aborted.

```
{  
    "testName": "Functional.bc",  
    "fail_case": [  

```

(continues on next page)

(continued from previous page)

```

    {
        "fail_regexp": "some test regexp",
        "fail_message": "some test message"
    },
    {
        "fail_regexp": "Bug",
        "fail_message": "Bug or Oops detected in system log",
        "use_syslog": 1
    }
],
"specs":
[
    {
        "name": "default",
        "EXPR": "3+3",
        "RESULT": "6"
    }
]
}

```

These variables are turned into environment variables by the overlay generator and are used with the function `fail_check_cases` which is called during the ‘post test’ phase of the test.

Note that the above items would be turned into the following environment variables internally in the fuego system:

- FUNCTIONAL_BC_FAIL_CASE_COUNT=2
- FUNCTIONAL_BC_FAIL_PATTERN_0="some test regexp"
- FUNCTIONAL_BC_FAIL_MESSAGE_0="some test message"
- FUNCTIONAL_BC_FAIL_PATTERN_1="Bug"
- FUNCTIONAL_BC_FAIL_MESSAGE_1="Bug or Oops detected in system log"
- FUNCTIONAL_BC_FAIL_1_SYSLOG=true

37.6 Catalog of current plans

Fuego, as of January 2017, only has a few testplans defined.

Here is the full list:

- testplan_default
- testplan_mmc
- testplan_sata
- testplan_usbstor
- testplan_bc_add
- testplan_bc_mult
- testplan_hello_world_fail
- testplan_hello_world_random

The storage-related testplans (mmc, sata, and usbstor) allow the test spec to configure the appropriate following variables:

- MOUNT_BLOCKDEV
- MOUNT_POINT
- TIMELIMIT
- NPROCS

Both the ‘bc’ and ‘hello_world’ testplans are example testplans to demonstrate how the testplan system works.

The ‘bc’ testplans are for selecting different operations to test in ‘bc’. The ‘hello_world’ testplans are for selecting different results to test in ‘hello_world’

38.1 PROGRAM

`parser.py`

38.2 DESCRIPTION

`parser.py` is a Python program that is used by each test to parse the test log for a test run, check the threshold(s) for success or failure, and store the data used to generate charts.

Each benchmark should include an executable file called ‘`parser.py`’ in its test home directory (`/fuego-core/tests/Benchmark.<testname>`). Functional tests may also provide a `parser.py`, when they return more than a single testcase result from the test. However, this is optional. If a Functional test does not have a `parser.py` script, then a generic one is used (called `generic_parser.py`), that just sets the result for the test based on the return code from the test program and the single result from running `log_compare` in the `test_processing` portion of the test script.

The overall operation of `parser.py` is as follows: `parser.py` reads the test log for the current run and parses it, extracting one or more testcase or measure results. These are stored in a python dictionary which is passed to the results processing engine. Normally the parsing is done by scanning the log using simple regular expressions. However, since this is a python program, an arbitrarily complex parser can be written to extract the result data from the test log.

38.2.1 Outline

The program usually has the following steps:

- Import the parser library
- Specify a search pattern for finding one or more measurements (or testcases) from the test log
- Call the `parse_log` function, to get a list of matches for the search pattern
- Build a dictionary of result values
- Call the `process` function.
 - The `process` function evaluates the results from the test, and determines the overall pass/fail status of a test, based on a *criteria.json* file
 - The `process` function also saves the information to the aggregate results file for this test (`flat_plot_data.txt`), and re-generates the chart data for the test (`flot_chart_data.json`).

38.2.2 Testcase and measure names

The `parser.py` program provides the name for the measures and testcases read from the test log file. It also provides the result values for these items, and passes the parsed data values to the processing routine.

These test names must be consistent in the `parser.py` program, `reference.json` file and the `criteria.json` file.

Please see Fuego naming rules for rules and guidelines for test names in the Fuego system.

38.3 SAMPLES

Here is a sample `parser.py` that does simple processing of a single metric. This is for `Benchmark.Dhrystone`.

Note the two calls to parser library functions: `parse_log()` and `process()`.

```
#!/usr/bin/python

import os, re, sys

sys.path.insert(0, os.environ['FUEGO_CORE'] + '/scripts/parser')
import common as plib

regex_string = "^ (Dhrystones.per.Second:)(\s*)([\d]{1,8}.?[\d]{1,3})(.*)$"

measurements = {}
matches = plib.parse_log(regex_string)

if matches:
    measurements['default.Dhrystone'] = [{"name": "Score", "measure": float(matches[0][2])}]

sys.exit(plib.process(measurements))
```

38.4 ENVIRONMENT and ARGUMENTS

`parser.py` uses the following environment variable:

- `FUEGO_CORE`

This is used to add `/fuego-core/scripts/parser` to the python system path, for importing the `common.py` module (usually as internal module name 'plib').

The parser library expects the following environment variables to be set:

- `FUEGO_RW`
- `FUEGO_RO`
- `FUEGO_CORE`
- `NODE_NAME`
- `TESTDIR`
- `TESTSPEC`
- `BUILD_NUMBER`

- BUILD_ID
- BUILD_TIMESTAMP
- PLATFORM
- FWVER
- LOGDIR
- FUEGO_START_TIME
- FUEGO_HOST
- Reboot
- Rebuild
- Target_PreCleanup
- WORKSPACE
- JOB_NAME

`parser.py` is called with the following invocation, from `function_processing`:

```
run_python $PYTHON_ARGS $FUEGO_CORE/tests/${TESTDIR}/parser.py
```

38.5 SOURCE

Located in `fuego-core/tests/${TESTDIR}/parser.py`.

38.6 SEE ALSO

- `parser_func_parse_log`, `parser_func_process`
- `function_processing`, Parser module API, `Benchmark_parser_notes`.

CRITERIA.JSON

39.1 Introduction

The `criteria.json` file is used to specify the criteria used to determine whether a test has passed or failed. The data and directives in a `criteria.json` file (referred to as the criteria data) allow Fuego to interpret test results and indicate ultimate success or failure of a test.

A test usually produces a number of individual testcase results or measurement values from the execution of the test. For functional tests, the criteria data can include the number of testcases in the test that must “PASS” or that may be allowed to “FAIL”. Or the criteria data can indicate specific testcase results that should be ignored. For benchmark tests, the criteria data includes threshold values for measurements taken by the benchmark, as well as operations (e.g. ‘less than’ or ‘greater than’), to use to determine if the value of a measurement should be interpreted as a “PASS” or a “FAIL”. Fuego uses the results of the test along with the criteria data, to determine the final top-level result of the test.

If no `criteria.json` file is provided, then a default is constructed based on the test results, consisting of the following:

```
{
  'tguid': <test_set_name>
  'max_fail': 0
}
```

39.1.1 Types of tests and pass criteria

A simple functional test runs a short sequence of tests, and if any one of them fails, then the test is reported as a failure. Since this corresponds to the default `criteria.json`, then most simple Functional tests do not need to provide a `criteria.json` file.

A complex functional test (such as LTP or glib) has hundreds or possibly thousands of individual test cases. Such tests often have some number of individual test cases that fail, but which may be safely ignored (either temporarily or permanently). For example, some test cases may fail sporadically due to problems with the test infrastructure or environment. Other tests may fail due to configuration choices for the software on the board. (For example, a choice of kernel config may cause some tests to fail - but this is expected and these fail results should be ignored).

Functional tests that are complex require a `criteria.json` file, to avoid failing the entire test because of individual testcases that should be ignored.

Finally, a Benchmark test is one that produces one or more “measurements”, which are test results with numeric values. In order to determine whether a result indicates a PASS or a FAIL result, Fuego needs to compare the numeric result with some threshold value. The `criteria.json` file holds the threshold value and operator used for making this comparison.

Different boards, or boards with different software installations or configurations, may require different pass criteria for the same tests. Therefore, the pass criteria are broken out into a separate file that can be adjusted at each test site,

and for each board. Ultimately, we would like testers to be able to share their pass criteria, so that each Fuego user does not have to determine these on their own.

39.2 Evaluation criteria

The criteria file lists “pass criteria” for test suites, test sets, test cases and measures. A single file may list one or more pass criteria for the test.

The criteria file may include count-based pass criteria, specific testcase lists, and measure reference values (thresholds).

The criteria file specifies the pass criteria for one or more test element results, by specifying the element’s test id (or tguid), and the criterion used to evaluate that element. Some results elements, such as test sets, are aggregates of other elements. For these, the criteria specify attributes of their child elements (like required counts, or listing individual children that must pass or fail).

The criteria file consists of a list of criterion objects (JSON objects), each of which specifies the tguid for the result element of the test, and additional data used to evaluate that element. tguids are generated by Fuego during the processing phase, and consist of statically defined strings unique to each test. You should look at a test’s *run.json* file to see the test element names for a test.

Here are the different operations that can be used for criteria:

- **max_fail** - specifies the maximum number of child elements that can fail, before causing this element to fail
 - by default, every aggregate element must have all it’s children pass, in order for it to pass (corresponding to a ‘max_fail’ of 0)
- **min_pass** - specifies the minimum number of child elements that must pass, in order for this element to pass
- **must_pass_list** - specifies a list of child elements, by name, that must pass for this element to pass
- **fail_ok_list** - specifies a list of child elements, by name, that may fail, without causing this element to fail
- **reference** - specifies a reference value used as a threshold to evaluate where a number value for this element represents pass or fail.
 - the reference object has two sub-attributes:
 - * **value** - the reference value (threshold)
 - * **operator** - the test between the result and the reference value

The operator can be one of the following strings:

- **gt** - result must be greater than the reference value
- **ge** - result must be greater than or equal to the reference value
- **lt** - result must be less than the reference value
- **le** - result must be less than or equal to the reference value
- **eq** - result must equal the reference value
- **ne** - result must not equal the reference value
- **bt** - result is between two reference values (or equal to one of them)

In case the reference object has an operator of ‘bt’, the ‘value’ field should have a string consisting of two numbers separated by a ‘,’. For example, to indicate that the result value should be between 4 and 5, the ‘value’ field should have the string “4,5”. Note that the comparison for ‘between’ also succeeds for equality. So in the example case of a reference value of “4,5”, the test would pass if the test result was exactly 4, or exactly 5, or any number between 4 and 5.

Note: The equality and inequality operators ('eq' and 'ne') are less likely to be useful for numerical evaluations of most benchmark measures, but are provided for completeness. These are useful if a test reports numerical results from within a small set of numbers (like 0 and 1).

39.3 Customizing the criteria.json file for a board

A Fuego user can customize the pass criteria for a board, by making a copy of the `criteria.json` file, manually editing the contents, and putting it in a specific directory with a specific filename, so Fuego can find it.

39.3.1 Using an environment variable

A Fuego user can specify their own path to the criteria file to use for a test using the environment variable `FUEGO_CRITERIA_JSON_PATH`. This can be set in the environment variables block in the Jenkins job for a test, if running the Fuego test from Jenkins, or in the shell environment prior to running a Fuego test using 'ftc'.

For example, the user could do the following:

- `$ export FUEGO_CRITERIA_JSON_PATH=/tmp/my-criteria.json`
- `$ ftc run-test -b board1 -t Functional.foo`

39.3.2 Using a board-specific directory

More commonly, a user can specify a board-specific criteria file, by placing the file under either `/fuego-rw/boards` or `/fuego-ro/boards`

When Fuego does test evaluation, it searches for the the criteria file to use, by looking for the following files in the indicated order:

- `$FUEGO_CRITERIA_JSON_PATH`
- `/fuego-ro/boards/{board}-{testname}-criteria.json`
- `/fuego-rw/boards/{board}-{testname}-criteria.json`
- `/fuego-core/tests/{testname}/criteria.json`

As an example, a user could customize the criteria file as follows:

- `$ cp /fuego-core/tests/Benchmark.Dhrystone/criteria.json /fuego-rw/boards/board1-Benchmark.Dhrystone-criteria.json`
- `$ edit /fuego-rw/boards/board1-Benchmark.Dhrystone-criteria.json`
 - Alter the reference value for the tguid 'default.Dhrystone.Score' to reflect a value appropriate for their board ('board1' in this example)
- (execute the job 'board1.default.Benchmark.Dhrystone' in Jenkins)
 - Fuego will use the criteria file for board1 in `/fuego-rw` instead of the default `criteria.json` file in the test's home directory

39.4 Examples

Here are some example `criteria.json` files:

39.4.1 Benchmark.dbench

```
{
  "schema_version": "1.0",
  "criteria": [
    {
      "tguid": "default.dbench.Throughput",
      "reference": {
        "value": 100,
        "operator": "gt"
      }
    },
    {
      "tguid": "default.dbench",
      "min_pass": 1
    }
  ]
}
```

The interpretation of this criteria file is that the measured value of `dbench.Throughput` (the result value) must have a value greater than 100. Also, at least 1 measure under the `default.dbench` test must pass, for the the entire test to pass.

39.4.2 Simple count

```
{
  "schema_version": "1.0",
  "criteria": [
    {
      "tguid": "default",
      "max_fail": 2
    }
  ]
}
```

The interpretation of this criteria file is that the test may fail up to 2 individual test cases, under the `default` test set, and still pass.

39.4.3 Child results

```
{
  "schema_version": "1.0",
  "criteria": [
    {
      "tguid": "syscall",
      "min_pass": 1000,
      "max_fail": 5
    },
    {
      "tguid": "timers",
      "fail_ok_list": ["leapsec_timer"]
    },
    {
      "tguid": "pty",
      "must_pass_list": ["hangup01"]
    }
  ]
}
```

The interpretation of this criteria file is that, within the `syscall` test set, a minimum of 1000 testcases must pass, and no more than 5 fail, in order for that set to pass. Also, in the test set `timers`, if the testcase `leapsec_timer` fails, it will not cause the entire test to fail. However, in the test set `pty`, the testcase `hangup01` must pass for the entire test to pass.

39.5 Schema

The schema for the `criteria.json` file is contained in the `fuego-core` repository at: `scripts/parser/fuego-criteria-schema.json`.

Here it is (as of Fuego 1.2):

```
{
  "$schema": "http://json-schema.org/schema#",
  "id": "http://www.fugetest.org/download/fuego_criteria_schema_v1.0.json",
  "title": "criteria",
  "description": "Pass criteria for a test suite",
  "definitions": {
    "criterion": {
      "title": "criterion ",
      "description": "Criterion for deciding if a test (test_set, test_case or ↪measure) passes",
      "type": "object",
      "properties": {
        "tguid": {
          "type": "string",
          "description": "unique identifier of a test (e.g.: Sequential_Output. ↪CPU)"
        },
        "min_pass": {
          "type": "number",

```

(continues on next page)

(continued from previous page)

```

        "description": "Minimum number of tests that must pass"
    },
    "max_fail": {
        "type": "number",
        "description": "Maximum number of tests that can fail"
    },
    "must_pass_list": {
        "type": "array",
        "description": "Detailed list of tests that must pass",
        "items": {
            "type": "string"
        }
    },
    "fail_ok_list": {
        "type": "array",
        "description": "Detailed list of tests that can fail",
        "items": {
            "type": "string"
        }
    },
    "reference": {
        "type": "object",
        "description": "Reference measure that is compared to a result_
↪measure to decide the status",
        "properties": {
            "value": {
                "type": [
                    "string",
                    "number",
                    "integer"
                ],
                "description": "A value (often a threshold) to compare_
↪against. May be two numbers separated by a comma for the 'bt' operator."
            },
            "operator": {
                "type": "string",
                "description": "Type of operation to compare against",
                "enum": [
                    "eq",
                    "ne",
                    "gt",
                    "ge",
                    "lt",
                    "le",
                    "bt"
                ]
            }
        }
    },
    "required": [
        "value",
        "operator"
    ]
}

```

(continues on next page)

(continued from previous page)

```

        }
    },
    "required": [
        "tguid"
    ]
},
{
    "type": "object",
    "properties": {
        "schema_version": {
            "type": "string",
            "description": "The version number of this JSON schema",
            "enum": [
                "1.0"
            ]
        },
        "criteria": {
            "type": "array",
            "description": "A list of criterion items",
            "items": {
                "$ref": "#/definitions/criterion"
            }
        }
    },
    "required": [
        "schema_version",
        "criteria"
    ]
}

```

39.6 Compatibility with previous Fuego versions

The `criteria.json` file replaces the `reference.log` file that was used in versions of Fuego prior to 1.2. If a test is missing a `criteria.json` file, and has a `reference.log` file, then Fuego will read the `reference.log` file and use its data as the pass criteria for the test.

Previously, Fuego (and its predecessor JTA) supported pass criteria functionality in two different ways:

- Functional test pass/fail counts
- Benchmark measure evaluations

39.6.1 Functional test pass/fail counts

For functional tests counts of positive and negative results were either hard-coded into the base scripts for the test, as arguments to the `log_compare()` in each test's `test_processing()` function, or they were specified as variables, read from the board file, and applied in the `test_processing()` function.

For example, the `Functional.OpenSSL` test used values of 176 pass and 86 fails (see `fuego-core/tests/Functional.OpenSSL/openssl.sh` in `fuego-1.1`) to evaluate the result of this test.

```
log_compare "$TESTDIR" "176" "${P_CRIT}" "p"
log_compare "$TESTDIR" "86" "${N_CRIT}" "n"
```

But tests in JTA, such as `Functional.LTP.Open_Posix` expected the variables `LTP_OPEN_POSIX_SUBTEST_COUNT_POS` and `LTP_OPEN_POSIX_SUBTEST_COUNT_NEG` to be defined in a the board file for the device under test.

For example, the board file might have lines like the following:

```
LTP_OPEN_POSIX_SUBTEST_COUNT_POS="1232"
LTP_OPEN_POSIX_SUBTEST_COUNT_NEG="158"
```

These were used in the `log_compare` function of the base script of the test like so:

```
log_compare "$TESTDIR" $LTP_OPEN_POSIX_SUBTEST_COUNT_POS "${P_CRIT}" "p"
log_compare "$TESTDIR" $LTP_OPEN_POSIX_SUBTEST_COUNT_NEG "${N_CRIT}" "n"
```

Starting with Fuego version 1.2, these would be replaced with `criteria.json` files like the following:

For `Functional.OpenSSL`:

```
{
  "schema_version": "1.0",
  "criteria": [
    {
      "tguid": "OpenSSL",
      "min_pass": 176,
      "max_fail": 86
    }
  ]
}
```

For `Functional.LTP.Open_Posix`:

```
{
  "schema_version": "1.0",
  "criteria": [
    {
      "tguid": "LTP.Open_Posix",
      "min_pass": 1232,
      "max_fail": 158
    }
  ]
}
```

FIXTHIS - should there be 'default' somewhere in the preceding tguids?

39.6.2 Benchmark measure evaluations

For Benchmark programs, the pass criteria consists of one or more measurement thresholds that are compared with the results produced by the Benchmark, along with the operator to be used for the comparison.

In JTA and Fuego 1.1 this data was contained in the `reference.log` file.

TOOLS.SH

The `tools.sh` file provides the definitions for variables used for each platform's toolchains.

A single `tools.sh` file is located at the directory `/fuego-ro/toolchains/tools.sh`. This file uses the `PLATFORM` environment variable to load a `tools.sh` file for a particular toolchain.

An individual platform tools file should be named `<PLATFORM>-tools.sh`, and be located in the `/fuego-ro/toolchains` directory.

A new `<PLATFORM>-tools.sh` files must be added whenever a toolchain or SDK is added to the system.

For example, for the platform `poky-qemuarm`, the file `/fuego-ro/toolchains/poky-qemuarm-tools.sh` sets the variables needed to compile programs with that toolchain.

The variables that should be exported are:

- `CC` - C compiler
- `CXX` - C++ compiler
- `CPP` - C pre-processor
- `CXXCPP` - C++ pre-processor
- `CONFIGURE_FLAGS` - flags for the configure script
- `RANLIB` - archive index generator (for libs)
- `AS` - assembler
- `LD` - linker
- `ARCH` - architecture
- `CROSS_COMPILE` - tool prefix used to build the kernel
- `PREFIX` - prefix used with most tools
- `HOST` - used with `configure --host=$HOST`, to specify the machine you are building for
- `SDKROOT` - used as prefix for `/usr/lib` and `/usr/include` directories

The above variables are directly referenced by the Fuego system.

A few other variables may be used optionally by the build instructions for individual tests.

- `CFLAGS`
- `LDLFLAGS`

40.1 Variable usage details

Note: Note that some tools variables are referenced in patch files. These don't count as uses from `-tools.sh`, because they are defined as part of the program build instructions with the program itself.

Here are some specific tools variables and what tests use them:

- CFLAGS - compiler flags
 - used by Benchmark.netpipe, Benchmark.cyclictest, Benchmark.tiobench, Benchmark.dbench, Benchmark.ffsh, Benchmark.Dhrystone, Benchmark.lmbench2, Benchmark.himeno, Benchmark.nbench_byte, Benchmark.linpack, Benchmark.GLMark, Benchmark.Whetstone, Functionall.synctest, Functiona.posixtestsuite, Functiona.scrashme, LTP, Functional.rmaptest, Functional.linus_stress, Functional.crashme
- LDFLAGS - linker flags
 - used by Benchmark.netpipe, Benchmark.cyclictest, Benchmark.Dhrystone, Benchmark.signaltest, Benchmark.Whetstone, Functional.synctest, Functiona.posixtestsuite, Functiona.scrashme, LTP, Functional.rmaptest, Functional.linus_stress, Functional.crashme
- HOST - this is passed to configure with `--host=$HOST`
 - used by Benchmark.aim7, Benchmark.bonnie, Benchmark.dbench, Benchmark.ffsb, Benchmark.x11perf, Benchmark.ipperf, Benchmark.gtkperf, Functional.ft2demos, netperf, Functional.glib, and Functional.stress.
- SDKROOT - used as prefix for `/usr/include` and `/usr/lib` directories and files in builds
 - used by Benchmark.aim7, Benchmark.blobsallad, Benchmark.GLMark, Benchmark.GLMark, Benchmark.GLMark, Benchmark.gtkperf, Functional.aiostress, Functional.zlib, LTP, and Functional.ft2demos

CODING STYLE

This page described the coding style conventions used in Fuego.

Please adhere to these conventions, so that the code has a more uniform style and it is easier to maintain. Not all code in Fuego adheres to these styles. As we work on code, we will convert it to the preferred style over time. New code should adhere to the preferred style.

Fuego code consists mostly of shell script and python code.

41.1 Indentation and line length

We prefer indentation to be 4 spaces, with no tabs.

It is preferred to keep lines within 80 columns. However, this is not strict. If a string constant causes a line to run over 80 columns, that is OK.

Some command sequences passed to the ‘report’ function may be quite long and require that they be expressed on a single line. In that case, you can break them up onto multiple lines using shell continuation lines.

41.2 Trailing whitespace

Lines should not end in trailing whitespace. That is: ‘grep ”\$” *’ should always be empty.

You can do this with: ‘grep -R ”\$” *’ in the directory you’re working in, and fix the lines manually.

Or, another method, if you’re using vim, is to add an autocmd to your .vimrc to automatically remove whitespace from lines that you edit.

This line in your ~/.vimrc:

```
autocmd FileType sh,c,python autocmd BufWritePre <buffer> %s/\s\+$//e
```

automatically removes whitespace from all line endings in shell, C, and python files that are saved from vim.

Or, a third method of dealing with this automatically is to have git check for whitespace errors using a configuration option, or a hook. See

<https://stackoverflow.com/questions/591923/make-git-automatically-remove-trailing-whitespace-before-committing#592014>

Also, script files should not end in blank lines.

41.3 Shell features

Shell scripts which run on the device-under-test (DUT or board), SHOULD restrict themselves to POSIX shell features only. Do not assume you have any shell features on the target board outside of those supported by ‘busybox ash’.

Try running the program `checkbashisms` on your target-side code, to check for any non-POSIX constructs in the code.

The code in `fuego_test.sh` is guaranteed to run in bash, and may contain bashisms, if needed. If equivalent functionality is available using POSIX features, please use those instead. Please avoid esoteric or little-known bash features. (Or, if you use such features, please comment them.)

Another useful tool for checking your shell code is a program called ‘ShellCheck’. See <https://github.com/koalaman/shellcheck>. Most distributions have a package for `shellcheck`.

There are a few conventions for avoiding using too many external commands in shell scripts that execute on the DUT. To check for a process, use `ps` and `grep`, but to avoid having `grep` find itself, use a wildcard in the search pattern. Like so: ‘`ps | grep [f]oo`’ (rather than ‘`ps | grep foo | grep -v grep`’).

41.4 Python style

Python code (such as parser code, the overlay generator, `ftc` and other helper scripts), should be compliant with <https://www.python.org/dev/peps/pep-0008/>. As with shell code, there is a lot of legacy code in Fuego that is not currently compliant with PEP 8. We will convert legacy code to the correct style as changes are made over time.

Here are a few more conventions for Fuego code:

- Strings consisting of a single character should be declared use single-quotes
- Strings consisting of multiple characters should declared using double-quotes, unless the string contains a double-quote. In that case, single-quotes should be using for quoting, to avoid having to escape the double-quote.

Note that there is a `fuego` lint test (selftest), called `Functional.fuego_lint`. It only checks a few files at the moment, but the plan is to expand it to check additional code in the future.

TESTPLAN REFERENCE

In Fuego, a testplan is used to specify a set of tests to execute, and the settings to use for each one.

42.1 Filename and location

A testplan is in json format, and can be located in two places:

- As a file located in the directory `fuego-core/overlay/testplans`.
- As a here document embedded in a batch test script (`fuego_test.sh`)

A testplan file name should start with the prefix “**testplan_**”, and end with the extension “.json”.

A testplan here document should be preceded by a line starting with `BATCH_TESTPLAN=` and followed by a line starting with “`END_TESTPLAN`”.

42.2 Top level attributes

The top level objects that may be defined in a testplan are:

- `testPlanName`
- `tests`
- `default_timeout`
- `default_spec`
- `default_reboot`
- `default_rebuild`
- `default_precleanup`
- `default_postcleanup`

Each of these attributes, except for ‘tests’ has a value that is a string. Here are their meanings and legal values:

The `testPlanName` is the name of this testplan. It must match the filename that holds this testplan (without the “**testplan_**” prefix or “.json” extension. This object is required.

‘tests’ is a list of tests that are part of this testplan. See below for a detailed description of the format of an element in the ‘tests’ list. This object is required.

42.2.1 Default test settings

The testplan may also include a set of default values for test settings. The test settings for which defaults may be specified are:

- timeout
- spec
- reboot
- rebuild
- precleanup
- postcleanup

These values are used if the related setting is not specified in the individual test definition.

For example, the testplan might define a `default_timeout` of “15m” (meaning 15 minutes). The plan could indicate timeouts different from this (say 5 minutes or 30 minutes) for individual tests, but if a test in the testplan doesn’t indicate its own timeout it would default to the one specified as the default at the top level of the testplan.

The ability to specify per-plan and per-test settings makes it easier to manage these settings to fit the needs of your Fuego board or lab.

Note that if neither the individual test nor the testplan provide a default value is not provided, then a Fuego global default value for that setting will be used.

Note that `default_spec` specifies the name of the test spec to use for the test (if one is not specified for the individual test definition). The name should match a spec that is defined for every test listed in the plan. Usually this will be something like “default”, but it could be something that is common for a set of tests, like ‘mmc’ or ‘usb’ for filesystem tests.

See the individual test definitions for descriptions of these different test settings objects.

42.3 Individual test definitions

The ‘tests’ object is a list of objects, each of which indicates a test that is part of the plan. The objects included in each list element are:

- testName
- spec
- timeout
- reboot
- rebuild
- precleanup
- postcleanup

All object values are strings.

42.3.1 TestName

The ‘testName’ object has the name of a Fuego test included in this plan. It should be the fully-qualified name of the test (that is, it should include the “Benchmark.” or “Functional.” prefix.) This attribute of the test element is required.

42.3.2 Spec

The ‘spec’ object has the name of the spec to use for this test. It should match the name of a valid test spec for this test. If ‘spec’ is not specified, then the value of “default_spec” for this testplan will be used.

42.3.3 Timeout

The timeout object has a string indicating the timeout to use for a test. The string is positive integer followed by a single-character units-suffix. The units suffixes available are:

- ‘s’ for seconds
- ‘m’ for minutes
- ‘h’ for hours
- ‘d’ for days

Most commonly, a number of minutes is specified, like so:

- “default_timeout” : “15m”,

If no ‘timeout’ is specified, then the value of ‘default_timeout’ for this testplan is used.

42.3.4 Reboot

The ‘reboot’ object has a string indicating whether to reboot the board prior to the test. It should have a string value of ‘true’ or ‘false’.

42.3.5 Rebuild

The ‘rebuild’ object has a string indicating whether to rebuild the test software, prior to executing the test. The object value must be a string of either ‘true’ or ‘false’.

If the value is ‘false’, then Fuego will do the following, when executing the test:

- If the test program is not built, then build it
- If the test program is already built, then use the existing test program

If the value is ‘true’, then Fuego will do the following:

- Remove any existing program build directory and assets
- Build the program (including fetching the source, unpacking it, and executing the instructions in the test’s “test_build” function)

42.3.6 Precleanup

The ‘precleanup’ flag indicates whether to remove all previous test materials on the target board, prior to deploying and executing the test. The object value must be a string of either ‘true’ or ‘false’.

42.3.7 Postcleanup

The ‘postcleanup’ flag indicates whether to remove all test materials on the target board, after the test is executed. The flag value must be a string of either ‘true’ or ‘false’.

42.4 Test setting precedence

Note that the test settings are used by the plan at job creation time, to set the command line arguments that will be passed to `ftc run-test` by the Jenkins job, when it is eventually run.

A user can always edit a Jenkins job (for a Fuego test), to override the test settings for that job.

The precedence of the settings encoded into the job definition at job creation time are:

- Testplan individual test setting (highest priority)
- Testplan default setting
- Fuego default setting

The precedence of settings at job execution time are:

- ‘ftc run-test’ command line option setting (highest priority)
- Fuego default setting

INDICES AND TABLES

- `genindex`
- `modindex`
- `search`

43.1 Sandbox

This page is for testing different elements of reStructuredText markup, in the Fuego documentation. This is intended to be used during the conversion from wiki pages, to make sure all important formatting is preserved.

43.2 Page Level Header (H1)

Here is some text

43.2.1 Chapter Level Header (H2)

This is the start of a level 2 section

Level 3 header

Does this actually do a level 3 header?

Level 4 header

Some content here at level 4

Level 5 header

Some content here at level 5

Some Example Markup

Here is a **bold** word, and an *italic* word. Also here is a `filename.txt`, and a `command -a arg`.

Here is a glossary terms: board

Admonition test

Reference test

test 1

Sandbox2_

Here is an attempt to refer to a page: ``Sandbox2`_`

Result:

- cover text of ``Sandbox2`_` - Fail, it includes the tics and underscore
- URL of `.../_build/html/Sandbox.html#id1` - FAIL, it's not to `Sandbox2.html`

test 2

Pointer to slandbox page1 `<sand1>_`

Here is another attemp to refer to a page. Try to refer to a section on a page, with cover text: `Pointer to slandbox page1`

Result:

- cover text of `Pointer to slandbox page1` - OK
- URL of `.../_build/html/sand1` - FAIL, it's not to `sand1.html`

test 3

Reference: `:ref:`Pointer to slandbox page3 <sndbx2>``

Here is another way to refer to a section on a page, with cover text: `Pointer to slandbox page3`

Result from default Ubuntu 16.03 Sphinx:

- cover text of `"Pointer to slandbox page3"`
- NO URL! - FAIL

Result from `python3 venv py3-sphinx` Sphinx installation:

- cover text of `"Pointer to slandbox page3"`
- URL of `.../_build/html/Sandbox2.html#sndbx2` - PASS

test 4

Reference: `:ref:`Pointer to slandbox2 page4 <Sandbox2>``

Here is another way to refer to a section on a page, with cover text: `Pointer to slandbox2 page4`

Result:

- cover text of `Pointer to slandbox2 page4`
- NO URL! - FAIL

test 5

Reference ``Pointer to slandbox2 page5 <Sandbox2.html>`_`

Here is another way to refer to a section on a page, with cover text: `Pointer to slandbox2 page5`

Result:

- cover text of `Pointer to slandbox page5`
- URL = `Sandbox2.html` - PASS

Conclusion. There doesn't seem to be a way to refer to a page or a section heading on a page, unless it is marked with an anchor. (unless you reference the page with it's .html extension)

test 6

Reference ``Pointer to slandbox2 test6 <chapheader2>``

Here is another way to refer to a section on a page, with cover text: `Pointer to slandbox2 test6`

Result:

- cover text of `Pointer to slandbox test6`
- URL = `.../_build/html/Sandbox2.html#chapheader2` - PASS

test 7

Reference ``Sandbox2``

Here is a way to refer to a whole page: `Sndbx2`

Result:

- cover text: `Sndbx2` - FAIL - this is the name in the first section on the page
- URL = `.../_build/html/Sandbox2.html` - PASS

test 8

Reference ``Sandbox2 <Sandbox2>``

Here is a way to refer to a whole page: [Sandbox2](#)

Result:

- cover text: Sandbox2 - PASS
- URL = `.../_build/html/Sandbox2.html` - PASS

:doc: items support cover text

test 9

Reference ``Cover text for Sandbox2 <Sandbox2>``

Here is a way to refer to a whole page: [Cover text for Sandbox2](#)

Result:

- cover text: Sandbox2 - PASS
- URL = `.../_build/html/Sandbox2.html` - PASS

:doc: items support cover text

test 10

Reference ``Cover text for Sandbox2 heading <test_header>``

Here is a way to refer to a whole page: Cover text for Sandbox2:test header

Result:

- cover text: “Cover text for Sandbox2:test header” - PASS
- URL = no link - FAIL

43.2.2 reference analysis

Sphinx is altogether too tricky when it comes to labels and cover text. It sometimes uses the section heading text for cover text for a label, rather than the label text itself, even when no cover text is specified.

This is true for :doc: items.

Which ways worked?

- Test 5 - but it's gross
- Test 6 - is ugly, for works for internal references
- Test 7 - works, but cover text is unpredictable
- Test 8,9 - are preferred for reference a whole page

43.2.3 Other tests

Stub heading

Toctree test

toctrees apparently refer to file (page) names. The items put into the tree are the section headings from those pages

Here's another toctree - this time with a caption

Following this is a hidden toctree

That's the end of the hiddent toctree.

I can keep doing this all day!!

Here's a literal block, with messed up indenting:

```
this is a test
this line should be next to that one
    this one should be indented
        even further indented
```

now we're done with the literal block

Notice there is no anchor or label before this section heading

These are items from the Sandbox2 page.

43.3 Sndbx2

This page header doesn't have the same name as the page.

This page is for testing different elements of reStructuredText markup, in the Fuego documentation. This is intended to be used during the conversion from wiki pages, to make sure all important formatting is preserved.

43.3.1 test header

Here is some material under the first header

43.4 Page Level Header2 (H1)

Here is some text

43.4.1 Chapter Level Header2

Stub header (h3)

Reference test

Test2 1

Reference: `Sandbox`_

Here is how you refer to a page: `**Sandbox`_**

Result: .../_build/html/Sandbox2.html#id4 - FAIL

Test2 2

Reference: `Sandbox2`_

Here is another way to refer to a page: `**Sandbox2`_**

Result: .../_build/html/Sandbox2.html#id6 - FAIL

Test2 3

Reference: Sandbox2_

Here is another way to refer to a page: **Sandbox2`_**

Result:

- cover = ???
- URL = .../_build/html/Sandbox2.html#id8 - FAIL

Test2 4

Description: anchor reference using underscore

Reference: `Pointer to slandbox page1 <sand1>`_

Here is one way to refer to a section on a page, with cover text: [Pointer to slandbox page1](#)

Result:

- cover - Pointer to slandbox page1
- URL - .../_build/html/sand1 - FAIL

Test2 5

Here is another way to refer to a section on a page, with cover text:

This uses a ref:

Reference: `:ref:`Pointer to slandbox page2 <sand1>``

Here is the ref: *Pointer to slandbox page2*

Result:

- cover = Pointer to slandbox page2
- URL = `.../_build/html/Sandbox.html#sand1` - PASS

Conclusion:

- you **MUST** have an anchor to use a ref
- only a `:ref:` gives you the page name (Sandbox.html) and the anchor ref (sand1)

Test2 more

Here is one way to refer to a section on a page, with cover text: This one uses trailing `_` and the page name: *Pointer to slandbox2 page*

Here is one way to refer to a section on a page, with cover text: This one uses an href and the page name: *Pointer to slandbox2 page*

Here is another way to refer to a section on a page, with cover text: This one uses a trailing `_` and the section name: *Pointer to slandbox2 page header*

Here is another way to refer to a section on a page, with cover text: This one uses a ref and the section name: *Pointer to slandbox2 page header*

Which ways worked?

43.5 cmd

43.5.1 NAME

cmd

43.5.2 SYNOPSIS

- `cmd <command and args>`

43.5.3 DESCRIPTION

The ‘cmd’ function is used to execute a command on the board.

It is important to quote arguments with spaces, so that they don’t get separated during execution.

It adds a statement to the devlog, and then calls `ov_transport_cmd`. It’s basically a simple wrapper around that function.

43.5.4 EXAMPLES

Here is an example

```
cmd "echo 'this echo is on the target'; uptime; $BOARD_TESTDIR/fuego.$TESTDIR/some_
↪program"
```

43.5.5 ENVIRONMENT and ARGUMENTS

The arguments are passed unchanged to `ov_transport_cmd`.

43.5.6 RETURN

Returns non-zero on error

43.5.7 SOURCE

Located in “scripts/functions.sh”

43.5.8 SEE ALSO

- `ov_transport_cmd`, `report_devlog`

43.6 run.json

43.6.1 Summary

The `run.json` file has data about a particular test run. It has information about the test, including the results for the test.

The format of portions of this file was inspired by the KernelCI API. See <https://api.kernelci.org/schema-test-case.html>

The results are included in an array of `test_set` objects, which can contain arrays of `test_case` objects, which themselves may contain measurement objects.

43.6.2 Field details

- **duration** - the amount of time, in milliseconds, that the test took to execute
 - If the test included a build, this time is included in this number
- **metadata** - various fields that are specific to Fuego
 - **attachments** - a list of the files that are available for this test - usually logs and such
 - **batch_id** - a string indicating the batch of tests this test was run in (if applicable)
 - **board** - the board the test was executed on
 - **build_number** - the Jenkins build number
 - **compiled_on** - indicates the location where the test was compiled
 - **fuego_core_version** - version of the fuego core system
 - **fuego_version** - version of the fuego container system
 - **host_name** - the host. If not configured, it may be 'local_host'
 - **job_name** - the Jenkins job name for this test run
 - **keep_log** - indicates whether the log is kept (???)
 - **kernel_version** - the version of the kernel running on the board
 - **reboot** - indicates whether a reboot was requested for this test run
 - **rebuild** - indicates whether it was requested to rebuild the source for this run
 - **start_time** - time when this test run was started (in seconds since Jan 1, 1970)
 - **target_postcleanup** - indicates whether cleanup of test materials on the board was requested for after test execution
 - **target_precleanup** - indicates whether cleanup of test materials on the board was requested for before test execution
 - **test_plan** - test plan being executed for this test run. May be 'None' if test was not executed in the context of a larger plan
 - **test_spec** - test spec used for this run
 - **testsuite_version** - version of the source program used for this run
 - * FIXTHIS - testsuite_version is not calculated properly yet
 - **timestamp** - time when this test run was started (in ISO 8601 format)
 - **toolchain** - the toolchains (or PLATFORM) used to build the test program
 - **workspace** - a directory on the host where test materials were extracted and built, for this test.
 - * This is the parent directory used, and not the specific directory used for this test.
- **name** - the name of the test
- **status** - the test result as a string. This can be one of:
 - PASS
 - FAIL
 - ERROR
 - SKIP

- **test_sets** - list of test_set objects, containing test results
- **test_cases** - list of test_case objects, containing test results
 - Each test_case object has:
 - * **name** - the test case name
 - * **status** - the result for that test case
- **measurements** - list of measurement objects, containing test results
 - For each measurement, the following attributes may be present:
 - * **name** - the measure name
 - * **status** - the pass/fail result for that test case
 - * **measure** - the numeric result for that test case

43.6.3 Examples

Here are some sample run.json files, from Fuego 1.2

Functional test results

This was generated using

```
ftc run-test -b docker -t Functional.hello_world
```

This example only has a single test_case.

```
{
  "duration_ms": 1245,
  "metadata": {
    "attachments": [
      {
        "name": "devlog",
        "path": "devlog.txt"
      },
      {
        "name": "devlog",
        "path": "devlog.txt"
      },
      {
        "name": "syslog.before",
        "path": "syslog.before.txt"
      },
      {
        "name": "syslog.after",
        "path": "syslog.after.txt"
      },
      {
        "name": "testlog",
        "path": "testlog.txt"
      },
    ]
  }
}
```

(continues on next page)

(continued from previous page)

```

    {
        "name": "consolelog",
        "path": "consolelog.txt"
    },
    {
        "name": "test_spec",
        "path": "spec.json"
    }
],
"board": "docker",
"build_number": "3",
"compiled_on": "docker",
"fuego_core_version": "v1.1-805adb0",
"fuego_version": "v1.1-5ad677b",
"host_name": "fake_host",
"job_name": "docker.default.Functional.hello_world",
"keep_log": true,
"kernel_version": "3.19.0-47-generic",
"reboot": "false",
"rebuild": "false",
"start_time": "1509662455755",
"target_postcleanup": true,
"target_precleanup": "true",
"test_plan": "None",
"test_spec": "default",
"testsuite_version": "v1.1-805adb0",
"timestamp": "2017-11-02T22:40:55+0000",
"toolchain": "x86_64",
"workspace": "/fuego-rw/buildzone"
},
"name": "Functional.hello_world",
"schema_version": "1.0",
"status": "PASS",
"test_sets": [
    {
        "name": "default",
        "status": "PASS",
        "test_cases": [
            {
                "name": "hello_world",
                "status": "PASS"
            }
        ]
    }
]
}
]
}

```

Benchmark results

Here is the `run.json` file for a run of the test `Benchmark.netperf` on the board `'ren1'` (which is a Renesas board in my lab).

```
{
  "duration_ms": 33915,
  "metadata": {
    "attachments": [
      {
        "name": "devlog",
        "path": "devlog.txt"
      },
      {
        "name": "devlog",
        "path": "devlog.txt"
      },
      {
        "name": "syslog.before",
        "path": "syslog.before.txt"
      },
      {
        "name": "syslog.after",
        "path": "syslog.after.txt"
      },
      {
        "name": "testlog",
        "path": "testlog.txt"
      },
      {
        "name": "consolelog",
        "path": "consolelog.txt"
      },
      {
        "name": "test_spec",
        "path": "spec.json"
      }
    ],
    "board": "ren1",
    "build_number": "3",
    "compiled_on": "docker",
    "fuego_core_version": "v1.2.0",
    "fuego_version": "v1.2.0",
    "host_name": "local_host",
    "job_name": "ren1.default.Benchmark.netperf",
    "keep_log": true,
    "kernel_version": "4.9.0-yocto-standard",
    "reboot": "false",
    "rebuild": "false",
    "start_time": "1509669904085",
    "target_postcleanup": true,
    "target_precleanup": "true",
    "test_plan": "None",
  }
}
```

(continues on next page)

(continued from previous page)

```

    "test_spec": "default",
    "testsuite_version": "v1.1-805adb0",
    "timestamp": "2017-11-03T00:45:04+0000",
    "toolchain": "poky-aarch64",
    "workspace": "/fuego-rw/buildzone"
  },
  "name": "Benchmark.netperf",
  "schema_version": "1.0",
  "status": "PASS",
  "test_sets": [
    {
      "name": "default",
      "status": "PASS",
      "test_cases": [
        {
          "measurements": [
            {
              "measure": 928.51,
              "name": "net",
              "status": "PASS"
            },
            {
              "measure": 59.43,
              "name": "cpu",
              "status": "PASS"
            }
          ],
          "name": "MIGRATED_TCP_STREAM",
          "status": "PASS"
        },
        {
          "measurements": [
            {
              "measure": 934.1,
              "name": "net",
              "status": "PASS"
            },
            {
              "measure": 56.61,
              "name": "cpu",
              "status": "PASS"
            }
          ],
          "name": "MIGRATED_TCP_MAERTS",
          "status": "PASS"
        }
      ]
    }
  ]
}

```

43.7 Fuego Test System

43.7.1 Welcome to Fuego!

Fuego is a test system specifically designed for embedded Linux testing. It supports automated testing of embedded targets from a host system, as its primary method of test execution.

Fuego consists of a host/target script engine, and over 100 pre-packages tests. These are installed in a docker container along with a Jenkins web interface and job control system, ready for out-of-the-box Continuous Integration testing of your embedded Linux project.

The idea is that in the simplest case, you just add your board, select or install a toolchain, and go!

Introduction presentation

Tim Bird gave some talks introducing Fuego, at various conferences in 2016. The slides and a video are provided below, if you want to see an overview and introduction to Fuego.

The slides are here: [Introduction-to-Fuego-LCJ-2016.pdf](#), along with a [YouTube video](#). You can find more presentations about Fuego on our wiki at: <http://fuegotest.org/wiki/Presentations>.

43.7.2 Getting Started

There are a few different ways to get started with Fuego:

1. Use the *Fuego Quickstart Guide* to get Fuego up and running quickly.
2. Or go through our *Install and First Test* tutorial to install Fuego and run a test on a single “fake” board. This will give you an idea of basic Fuego operations, without having to configure Fuego for your own board
3. Work through the documentation for *Installation*

Where to download

Code for the test framework is available in 2 git repositories:

- <https://bitbucket.org/fuegotest/fuego/>
- <https://bitbucket.org/fuegotest/fuego-core/>

The fuego-core directory resides inside the fuego directory. But normally you do not clone that repository directly. It is cloned for you during the Fuego install process. See the *Fuego Quickstart Guide* or the *Installing Fuego* page for more information.

43.7.3 Documentation

See the index below for links to the major sections of the documentation for Fuego. The major sections are:

- *Tutorials*
- *Installation and Administration*
- *User Guides*
- *Developer Resources*
- *API Reference*

43.7.4 Resources

Mailing list

Fuego discussions are held on the fuego mailing list:

- <https://lists.linuxfoundation.org/mailman/listinfo/fuego>

Note that this is a new list (as of September 2016). Previously, discussions about Fuego (and its predecessor JTA) were held on the ltsi-dev mailing list:

- <https://lists.linuxfoundation.org/mailman/listinfo/ltsi-dev>

Presentations

A number of presentations have been given on the Fuego test framework, and related projects (such as its predecessor JTA, and a derivative project JTA-AGL).

See the [Presentations](#) page on the Fuego wiki for a list of presentations that you can read or view for more information about Fuego.

43.7.5 Vision

The purpose of Fuego is to bring the benefits of open source to the testing process.

It can be summed up like this:

Note: Do for testing what open source has done for coding

There are numerous aspects of testing that are still done in an ad-hoc and company-specific way. Although there are open source test frameworks (such as Jenkins or LAVA), and open source test programs (such as cyllictest, LTP, linuxbench, etc.), there are lots of aspects of Linux testing that are not shared.

The purpose of Fuego is to provide a test framework for testing embedded Linux, that is distributed and allows individuals and organizations to easily run their own tests, and at the same time allows people to share their tests and test results with each other.

Historically, test frameworks for embedded Linux have been difficult to set up, and difficult to extend. Many Linux test systems are not easily applied in cross or embedded environments. Some very full frameworks are either not viewed as processor-neutral, and are difficult to set up, or are targeted at running tests on a dedicated group of boards or devices.

The vision of open source in general is one of sharing source code and capabilities, to expand the benefits to all participants in the ecosystem. The best way to achieve this is to have mechanisms to easily use the system, and easily share enhancements to the system, so that all participants can use and build on each others efforts.

The goal of Fuego is to provide a framework that any group can install and use themselves, while supporting important features like cross-compilation, host/target test execution, and easy test administration. Test administration consists of starting tests (both manually and automatically), viewing test results, and detecting regressions. Ease of use is critical, to allow testers to use tests that are otherwise difficult to individually set up, configure, and interpret the results from. It is also important to make it very easy to share tests (scripts, configuration, results parsing, and regression detection methods).

Some secondary goals of this project are the ability for 3rd parties to initiate or schedule tests on our hardware, and the ability to share our test results with others.

The use of Jenkins as the core of the test framework already supports many of the primary and secondary goals. The purpose of this project is to augment the Jenkins system to support embedded configurations of Linux, and to provide a place for centralized sharing of test configurations and collateral.

There is no such thing as a “Linux Test distribution”. Fuego aims to be this. It intends to provide test programs, a system to build, deploy and run them, and tools to analyze, track, and visualize test results.

For more details about a high-level vision of open source testing, please see *OSS Test Vision*.

43.7.6 Other Resources

Historical information

<http://elinux.org/Fuego> has some historical information about Fuego.

Things to do

Looking for something to do on Fuego? See the Fuego wiki for a list of projects, at: [Fuego To Do List](#)

43.8 FUEGO BUILD FLAGS

This variable may be defined in a board file or in a tools file, and is used to specify build attributes for test programs built by Fuego for that board or toolchain (respectively)

Currently, this can only have the value: **no_static**

Other values may be supported (in a space-separated list) in the future.

By default, many Fuego tests will attempt to build static binaries, as these require less dependencies on the target. However, some toolchains do not support compiling static binaries. For such a toolchain, this flag should be used in the toolchain shell script, to indicate to the Fuego build system to build dynamic programs instead.

This flag can also be used in a board file, to indicate to the Fuego build system to build dynamic programs for that board (whether or not the toolchain supports static linking or not).

Example: put this line in a toolchain setup script, (`/fuego-ro/toolchains/$TOOLCHAIN-tools.sh`) or in your board configuration file (`/fuego-ro/boards/myboard.board`):

```
FUEGO_BUILD_FLAGS="no_static"
```

43.9 FUEGO DEBUG

Note: FUEGO_DEBUG is now deprecated. Please use the newer FUEGO_LOGLEVELS feature instead of this. As of Fuego version 1.4, FUEGO_DEBUG is still supported for backwards compatibility.

The environment variable FUEGO_DEBUG is used to control debug output during execution of a Fuego test.

If this variable is not set, no debugging messages (or less messages) are produced.

The variable is a bitmask. If it is defined at all, then the script system will produce shell trace messages as part of the test log.

The following bitmask values can be used to turn on debugging for different parts of the system:

- 1 = debug the main execution of test phases
- 2 = debug the parser
- 4 = debug the criteria processor
- 8 = debug the chart generator code

Combinations are allowed, but must be in decimal.

Example:

```
export FUEGO_DEBUG=15
```

This would turn on debug messages for all areas.

43.10 FUEGO LOGLEVELS

The environment variable FUEGO_LOGLEVELS is used to control message output (including debug messages) during execution of a Fuego test.

43.10.1 Introduction

The FUEGO_LOGLEVELS variable specifies a string containing a list of areas and log level combinations, separated by commas. The area and loglevel are joined by a colon.

Here is an example:

```
export FUEGO_LOGLEVELS="deploy:verbose,criteria:debug"
```

Note that a sample of this line is provided in the Jenkins job for every test. It is, by default, commented out. However, you can easily turn on FUEGO_LOGLEVELS by uncommenting this line. You can customize the log level to use for different execution areas by changing the value of the variable.

To change this line in a Jenkins job, select the job in the Jenkins interface, then select “Configure”, and edit the line in “Execute Shell - Command” box, in the “Build” section of the job configuration.

If the FUEGO_LOGLEVELS variable is not set, the default logging level for all areas of test execution is “info”.

43.10.2 Log levels

There are 5 logging levels available, and messages from Fuego are categorized into these 5 different levels:

- error
- warning
- info
- verbose
- debug

Specifying a particular level means that all messages above that level will be output. Messages at level ‘error’ are always shown, no matter what log level is specified.

43.10.3 Execution areas

Area names correspond to phases, and to sub-phases of the test execution stepx. The following area names are supported:

- pre_test
- pre_check
- build
- makepkg
- deploy
- snapshot
- run
- post_test
- processing
- parser
- criteria
- charting

43.10.4 Output functions

With this feature, 5 new functions have been added to the Fuego core. These functions may be used in your test shell script (`fuego_test.sh`), so that your output may be managed the same way that core output is managed.

The following functions are available:

- `dprint` - print output if the message level is 'debug'.
- `vprint` - print output if the message level is 'debug' or 'verbose'.
- `iprint` - print output if the message level is 'debug', 'verbose', or 'info'.
- `wprint` - print output if the message level is 'debug', 'verbose', 'info', or 'warning'.
- `eprint` - print output always (message level 'error')

Deprecated FUEGO_DEBUG

`FUEGO_LOGLEVELS` replaces the earlier `FUEGO_DEBUG` variable for controlling debug output of Fuego. However, as of Fuego version 1.4, `FUEGO_DEBUG` is still supported for backwards compatibility.