
mush Documentation

Release 2.1

Simplistix Ltd

December 14, 2015

1	How Mush works	3
1.1	Constructing runners	3
1.2	Configuring Resources	5
1.3	Labels	10
1.4	Context manager resources	12
1.5	Testing	13
1.6	Debugging	14
2	Example Usage	15
2.1	Encapsulating the re-usable parts of scripts	16
2.2	Writing the specific parts of your script	18
2.3	Assembling the components into a script	19
2.4	Testing	20
3	API Reference	21
4	Installation Instructions	25
5	Development	27
5.1	Setting up a virtualenv	27
5.2	Running the tests	27
5.3	Building the documentation	27
5.4	Making a release	28
6	Changes	29
6.1	2.1 (14 December 2015)	29
6.2	2.0 (11 December 2015)	29
6.3	1.3 (21 October 2015)	29
6.4	1.2 (11 December 2013)	29
6.5	1.1 (27 November 2013)	30
6.6	1.0 (29 October 2013)	30
7	License	31
8	Indices and tables	33
	Python Module Index	35

Mush is a light weight type-based dependency injection framework aimed at enabling the easy testing and re-use of chunks of code that make up scripts.

How Mush works

Note: This documentation explains how Mush works using fairly abstract examples. If you’d prefer more “real world” examples please see the [Example Usage](#) documentation.

1.1 Constructing runners

Mush works by assembling a number of callables into a *Runner*:

```
from mush import Runner

def func1():
    print('func1')

def func2():
    print('func2')

runner = Runner(func1, func2)
```

Once assembled, a runner can be called any number of times. Each time it is called, it will call each of its callables in turn:

```
>>> runner()
func1
func2
```

More callables can be added to a runner:

```
def func3():
    print('func3')

runner.add(func3)
```

If you want to add several callables in one go, you can use the runner’s *extend()* method:

```
def func4():
    print('func4')

def func5():
    print('func5')

runner.extend(func4, func5)
```

Now, when called, the runner will call all five functions:

```
>>> runner()
func1
func2
func3
func4
func5
```

Runners can also be added together to create a new runner:

```
runner1 = Runner(func1)
runner2 = Runner(func2)
runner3 = runner1 + runner2
```

This addition does not modify the existing runners, but does give the result you'd expect:

```
>>> runner1()
func1
>>> runner2()
func2
>>> runner3()
func1
func2
```

This can also be done by passing runners in when creating a new runner or calling the extend method on a runner, for example:

```
runner1 = Runner(func1)
runner2 = Runner(func2)
runner4_1 = Runner(runner1, runner2)
runner4_2 = Runner()
runner4_2.extend(runner1, runner2)
```

In both cases, the results are as you would expect:

```
>>> runner4_1()
func1
func2
>>> runner4_2()
func1
func2
```

Finally, runners can be cloned, providing a way to encapsulate commonly used base runners that can then be extended for each specific use case:

```
runner5 = runner3.clone()
runner5.add(func4)
```

The existing runner is not modified, while the new runner behaves as expected:

```
>>> runner3()
func1
func2
>>> runner5()
func1
func2
func4
```


1.2 Configuring Resources

Where Mush becomes useful is when the callables in a runner either produce or require objects of a certain type. Given the right configuration, Mush will wire these together enabling you to write easily testable and reusable callables that encapsulate specific pieces of functionality. This configuration is done either imperatively, declaratively or using a combination of the two styles as described in the sections below.

For the examples, we'll assume we have three types of resources:

```
class Apple:
    def __str__(self):
        return 'an apple'
    __repr__ = __str__

class Orange:
    def __str__(self):
        return 'an orange'
    __repr__ = __str__

class Juice:
    def __str__(self):
        return 'a refreshing fruit beverage'
    __repr__ = __str__
```

1.2.1 Specifying requirements

When callables take parameters, Mush can be configured to pass objects of the correct type that have been returned from previous callables in the runner. For example, consider the following functions:

```
def apple_tree():
    print('I made an apple')
    return Apple()

def magician(fruit):
    print('I turned {0} into an orange'.format(fruit))
    return Orange()

def juicer(fruit1, fruit2):
    print('I made juice out of {0} and {1}'.format(fruit1, fruit2))
```

The requirements are specified by passing the required type in the *requires* parameter when adding the callable to the runner using *add()*. If more complex requirements need to be specified, a *requires* instance can be passed which accepts both positional and keyword parameters that specify the types required by the callable being added:

```
from mush import Runner, requires

runner = Runner()
runner.add(apple_tree)
runner.add(magician, requires=Apple)
runner.add(juicer, requires(Apple, fruit2=Orange))
```

Calling this runner will now manage the resources, collecting them and passing them in as configured:

```
>>> runner()
I made an apple
I turned an apple into an orange
I made juice out of an apple and an orange
```

1.2.2 Optional requirements

It may be that, while a callable needs an object of a particular type, a default can be used if no such object is present. Runners can be configured to take this into account. Take the following function:

```
def greet(name='stranger'):  
    print('Hello ' + name + '!!')
```

If a name is not always be available, it can be added to a runner as follows:

```
from mush import Runner, optional  
  
runner = Runner()  
runner.add(greet, requires=optional(str))
```

Now, when this runner is called, the default will be used:

```
>>> runner()  
Hello stranger!
```

The same callable can be added to a runner where the required strings is available:

```
from mush import Runner, optional  
  
def my_name_is():  
    return 'Slim Shady'  
  
runner = Runner(my_name_is)  
runner.add(greet, requires=optional(str))
```

In this case, the string returned will be used:

```
>>> runner()  
Hello Slim Shady!
```

1.2.3 Using parts of a resource

Resources can have attributes or items that are directly required by callables. For example, consider these two functions that return such resources:

```
class Stuff(object):  
    fruit = 'apple'  
    tree = dict(fruit='pear')  
  
def some_attributes():  
    return Stuff()  
  
def some_items():  
    return dict(fruit='orange')
```

Also consider this function:

```
def pick(fruit1, fruit2, fruit3):  
    print('I picked {0}, {1} and {2}'.format(fruit1, fruit2, fruit3))
```

All three can be added to a runner such that mush will pass the correct parts of the returns resources through to the `pick()` function:

```
from mush import Runner, attr, item, requires

runner = Runner(some_attributes, some_items)
runner.add(pick, requires(fruit1=attr(Stuff, 'fruit'),
                           fruit2=item(dict, 'fruit'),
                           fruit3=item(attr(Stuff, 'tree'), 'fruit')))
```

So now we can pick fruit from some interesting places:

```
>>> runner()
I picked apple, orange and pear
```

The `pick()` function, however, remains usable and testable on its own:

```
>>> pick('apple', 'orange', 'pear')
I picked apple, orange and pear
```

1.2.4 Specifying returned resources

As seen above, Mush will track resources returned by callables based on the type of any object returned. This is usually what you want, but in some cases you may want to specify something different. For example, if you have a callable that returns a sequence of resources, this would be added to a runner as follows:

```
from mush import Runner, returns_sequence

def all_fruit():
    print('I made fruit')
    return Apple(), Orange()

runner = Runner()
runner.add(all_fruit, returns=returns_sequence())
runner.add(juicer, requires(Apple, Orange))
```

Now, the juicer will use all the fruit returned:

```
>>> runner()
I made fruit
I made juice out of an apple and an orange
```

In rarer circumstances, you may need to override the types returned by a callable. This can be done as follows:

```
from mush import Runner, returns

class Tomato:
    def __str__(self):
        return 'a tomato'

class Cucumber:
    def __str__(self):
        return 'a cucumber'

def vegetables():
    print('I made vegetables')
    return Tomato(), Cucumber()

runner = Runner()
```

```
runner.add(vegetables, returns=returns(Apple, Orange))
runner.add(juicer, requires(Apple, Orange))
```

Now, even when a callable requires fruit, we can force it to be happy with vegetables:

```
>>> runner()
I made vegetables
I made juice out of a tomato and a cucumber
```

The `returns` indicator can be used even if a single object is returned.

Another way that the type used to track the resource can be different from the type of the resource itself is if a callable returns a mapping and Mush is configured to use the types from that mapping:

```
from mush import Runner, returns_mapping

def desperation():
    print('I sold vegetables as fruit')
    return {Apple: Tomato(), Orange: Cucumber()}

runner = Runner()
runner.add(desperation, returns=returns_mapping())
runner.add(juicer, requires(Apple, Orange))
```

One again, we can happily make juice out of vegetables:

```
>>> runner()
I sold vegetables as fruit
I made juice out of a tomato and a cucumber
```

1.2.5 Named resources

Sometimes the types of resources are too common for them to uniquely identify a resource:

```
def age():
    return 37

def meaning():
    return 42
```

A callable such as the following cannot be configured to require the correct resource from these two functions by type alone:

```
def profound(age, it):
    print('by the age of %s I realised the meaning of life was %s' % (
        age, it
    ))
```

For these situations, Mush supports the ability to name a resource using a string:

```
runner = Runner()
runner.add(age, returns='age')
runner.add(meaning, returns='meaning')
runner.add(profound, requires('age', it='meaning'))
```

Anywhere that a type can be used, a string name can be used instead:

```
>>> runner()
by the age of 37 I realised the meaning of life was 42
```

1.2.6 Declarative configuration

While the imperative configuration used so far means that callables do not need to be modified, the helpers for specifying requirements and return types can also be used as decorators:

```
from mush import requires

def apple_tree():
    print('I made an apple')
    return Apple()

@requires(Apple)
@returns('citrus')
def magician(fruit):
    print('I turned {0} into an orange'.format(fruit))
    return Orange()

@requires(fruit1=Apple, fruit2='citrus')
def juicer(fruit1, fruit2):
    print('I made juice out of {0} and {1}'.format(fruit1, fruit2))
    return Juice()
```

These can now be combined into a runner and executed. The runner will extract the requirements stored by the decorator and will use them to map the parameters as appropriate:

```
>>> runner = Runner(apple_tree, magician, juicer)
>>> runner()
I made an apple
I turned an apple into an orange
I made juice out of an apple and an orange
a refreshing fruit beverage
```

1.2.7 Hybrid configuration

The two styles of configuration are entirely interchangeable, with declarative requirements being inspected whenever a callable is added to a runner, and imperative requirements being taken whenever they are passed via the `add()` method:

```
@requires(Juice)
def packager(juice):
    print('I put {0} in a bottle'.format(juice))

def orange_tree():
    print('I made an orange')
    return Orange()

trees = Runner(apple_tree)
trees.add(orange_tree, returns='citrus')
runner = trees.clone()
runner.extend(juicer, packager)
```

This runner now ends up with bottled juice:

```
>>> runner()
I made an apple
I made an orange
I made juice out of an apple and an orange
I put a refreshing fruit beverage in a bottle
```

It's useful to note that imperative configuration will be used in preference to declarative configuration where both are present:

```
runner = trees.clone()
runner.add(juicer, requires('citrus', Apple))
```

This runner will give us juice made in a different order:

```
>>> runner()
I made an apple
I made an orange
I made juice out of an orange and an apple
a refreshing fruit beverage
```

1.3 Labels

One of the motivating reasons for Mush to be created was the ability to insert callables at a point in a runner other than the end. This allows abstraction of common sequences of calls without the risks of extracting them into a base class.

The points at which more callables can be inserted are created by specifying a label when adding a callable to the runner. This marks the point at which that callable is included so that it can be retrieved and appended to later. As an example, consider a ring and some things that can be done to it:

```
class Ring:
    def __str__(self):
        return 'a ring'

def forge():
    return Ring()

def engrave(ring):
    print('engraving {}'.format(ring))
```

These might be added to a runner as follows:

```
from mush import Runner

runner = Runner()
runner.add(forge, label='forged')
runner.add(engrave, label='engraved', requires=Ring)
```

Now, suppose we want to polish the ring before it's engraved and then package it up when we're done:

```
def polish(ring):
    print('polishing {}'.format(ring))

def package(ring):
    print('packaging {}'.format(ring))
```

We can insert these callables into the runner at the right points as follows:

```
runner['forged'].add(polish, requires=Ring)
runner.add(package, requires=Ring)
```

This results in the desired call order:

```
>>> runner()
polishing a ring
```

```
engraving a ring
packaging a ring
```

Now, suppose we want to polish the ring again after it's been engraved. We can insert another callable at the appropriate point:

```
def more_polish(ring):
    print('polishing {0} again'.format(ring))

runner['engraved'].add(more_polish, requires=Ring)
```

Mush wish do the right thing when this runner is called:

```
>>> runner()
polishing a ring
engraving a ring
polishing a ring again
packaging a ring
```

When using labels, it's often good to be able to see exactly what is in a runner, what order it is in and where any labels point. For this reason, the representation of a runner gives all this information:

```
>>> runner
<Runner>
  <function forge ...> requires() returns_result_type()
  <function polish ...> requires(Ring) returns_result_type() <-- forged
  <function engrave ...> requires(Ring) returns_result_type()
  <function more_polish ...> requires(Ring) returns_result_type() <-- engraved
  <function package at ...> requires(Ring) returns_result_type()
</Runner>
```

As you can see above, when a callable is inserted at a label, the label moves to that callable. You may wish to keep track of the initial point that was labelled, so Mush supports multiple labels at each point:

```
>>> runner = Runner()
>>> point = runner.add(forge)
>>> point.add_label('before_polish')
>>> point.add_label('after_polish')
>>> runner
<Runner>
  <function forge ...> requires() returns_result_type() <-- after_polish, before_polish
</Runner>
```

Now, when you add to a specific label, only that label is moved:

```
>>> point = runner['after_polish']
>>> point.add(polish)
>>> runner
<Runner>
  <function forge ...> requires() returns_result_type() <-- before_polish
  <function polish ...> requires() returns_result_type() <-- after_polish
</Runner>
```

Of course, you can still add to the end of the runner:

```
>>> runner.add(package)
<mush.modifier.Modifier...>
>>> runner
<Runner>
  <function forge ...> requires() returns_result_type() <-- before_polish
```

```
<function polish ...> requires() returns_result_type() <-- after_polish
<function package ...> requires() returns_result_type()
</Runner>
```

However, the point modifier returned by getting a label from a runner will keep on moving the label as more callables are added using it:

```
>>> point.add(more_polish)
>>> runner
<Runner>
  <function forge ...> requires() returns_result_type() <-- before_polish
  <function polish ...> requires() returns_result_type()
  <function more_polish ...> requires() returns_result_type() <-- after_polish
  <function package ...> requires() returns_result_type()
</Runner>
```

1.4 Context manager resources

A frequent requirement when writing scripts is to make sure that when unexpected things happen they are logged, transactions are aborted, and other necessary cleanup is done. Mush supports this pattern by allowing context managers to be added as callables:

```
from mush import Runner, requires

class Transactions(object):

    def __enter__(self):
        print('starting transaction')

    def __exit__(self, type, obj, tb):
        if type:
            print(obj)
            print('aborting transaction')
        else:
            print('committing transaction')
        return True

def a_func():
    print('doing my thing')

def good_func():
    print('I have done my thing')

def bad_func():
    raise Exception("I don't want to do my thing")
```

The context manager is wrapped around all callables that are called after it:

```
>>> runner = Runner(Transactions, a_func, good_func)
>>> runner()
starting transaction
doing my thing
I have done my thing
committing transaction
```

This gives it a chance to clear up when things go wrong:


```
>>> runner = Runner(Transactions, a_func, bad_func)
>>> runner()
starting transaction
doing my thing
I don't want to do my thing
aborting transaction
```

1.5 Testing

Mush has a couple of features to help with automated testing of runners. For example, if you wanted to test a runner that got configuration by calling a remote web service:

```
@returns('config')
def load_config():
    return json.loads(urllib2.urlopen('...').read())

@requires(item('config', 'username'), item('config', 'password'))
def do_stuff(username, password):
    print('doing stuff as ' + username + ' with ' + password)

runner = Runner(load_config, do_stuff)
```

When testing this runner, we may want to inject a hard-coded config. This can be done by cloning the original runner and replacing the `load_config()` callable:

```
>>> def test_config():
...     return dict(username='test', password='pw')
>>> test_runner = runner.clone()
>>> test_runner.replace(load_config, test_config)
>>> test_runner()
doing stuff as test with pw
```

If you have a base runner such as this:

```
def load_config():
    return json.loads(urllib2.urlopen('...').read())

def finalise_things():
    print('all done')

base_runner = Runner()
point = base_runner.add(load_config, returns='config', label='config')
point.add_label('body')
base_runner.add(finalise_things, label='ending')
```

That runner might be used for a specific script as follows:

```
@requires(item('config', 'username'), item('config', 'password'))
def do_stuff(username, password):
    print('doing stuff as ' + username + ' with ' + password)

runner = base_runner.clone()
runner['body'].add(do_stuff)
```

To test this runner, we want to use a dummy configuration and not have any finalisation take place. This can be achieved as follows:

```
>>> test_runner = Runner()
>>> test_runner.add(test_config, returns='config')
<mush.modifier.Modifier ...>
>>> test_runner += runner.clone(start_label='config', end_label='ending')
>>> test_runner()
doing stuff as test with pw
```

1.6 Debugging

Mush has a couple of features to aid debugging of runners. The first of which is that the representation of a runner will show everything in it, in the order it will be called and what each callable has been declared as requiring and returning along with where any labels currently point.

For example, consider this runner:

```
from mush import Runner

@returns('config')
def make_config():
    return {'foo': 'bar'}

@requires(item('config', 'foo'))
def connect(foo):
    return 'connection'

def process(connection):
    print('using ' + repr(connection))

runner = Runner()
point = runner.add(make_config, label='config')
point.add(connect)
runner.add(process)
```

To see how the configuration panned out, we would look at the `repr()`:

```
>>> runner
<Runner>
  <function make_config ...> requires() returns('config')
  <function connect ...> requires('config'['foo']) returns_result_type() <-- config
  <function process ...> requires() returns_result_type()
</Runner>
```

As you can see, there is a problem with this configuration that will be exposed when it is run. To help make sense of these kinds of problems, Mush will add more context when a `TypeError` or `ContextError` is raised.

Example Usage

To show how Mush works from a more practical point of view, let's start by looking at a simple script that covers several common patterns:

```
from argparse import ArgumentParser
from .configparser import RawConfigParser
import logging, os, sqlite3, sys

log = logging.getLogger()

def main():
    parser = ArgumentParser()
    parser.add_argument('config', help='Path to .ini file')
    parser.add_argument('--quiet', action='store_true',
                        help='Log less to the console')
    parser.add_argument('--verbose', action='store_true',
                        help='Log more to the console')
    parser.add_argument('path', help='Path to the file to process')

    args = parser.parse_args()

    config = RawConfigParser()
    config.read(args.config)

    handler = logging.FileHandler(config.get('main', 'log'))
    handler.setLevel(logging.DEBUG)
    log.addHandler(handler)
    log.setLevel(logging.DEBUG)

    if not args.quiet:
        handler = logging.StreamHandler(sys.stderr)
        handler.setLevel(logging.DEBUG if args.verbose else logging.INFO)
        log.addHandler(handler)

    conn = sqlite3.connect(config.get('main', 'db'))

    try:
        filename = os.path.basename(args.path)
        with open(args.path) as source:
            conn.execute('insert into notes values (?, ?)',
                         (filename, source.read()))
        conn.commit()
        log.info('Successfully added %r', filename)
    except:
```

```
log.exception('Something went wrong')

if __name__ == '__main__':
    main()
```

As you can see, the script above takes some command line arguments, loads some configuration from a file, sets up log handling and then loads a file into a database. While simple and effective, this script is hard to test. Even using the `TempDirectory` and `Replacer` helpers from the `TestFixtures` package, the only way to do so is to write one or more high level tests such as the following:

```
from testfixtures import TempDirectory, Replacer, OutputCapture
import sqlite3

class Tests(TestCase):

    def test_main(self):
        with TempDirectory() as d:
            # setup db
            db_path = d.getpath('sqlite.db')
            conn = sqlite3.connect(db_path)
            conn.execute('create table notes (filename varchar, text varchar)')
            conn.commit()
            # setup config
            config = d.write('config.ini', '''

[main]
db = %s
log = %s
''' % (db_path, d.getpath('script.log')), 'ascii')
            # setup file to read
            source = d.write('test.txt', 'some text', 'ascii')
            with Replacer() as r:
                r.replace('sys.argv', ['script.py', config, source, '--quiet'])
                main()
            # check results
            self.assertEqual(
                conn.execute('select * from notes').fetchall(),
                [('test.txt', 'some text')]
            )
```

The problem is that, in order to test the different paths through the small piece of logic at the end of the script, we have to work around all the set up and handling done by the rest of the script.

This also makes it hard to re-use parts of the script. It's common for a project to have several scripts, all of which get some config from the same file, have the same logging options and often use the same database connection.

2.1 Encapsulating the re-usable parts of scripts

So, let's start by looking at how these common sections of code can be extracted into re-usable functions that can be assembled by Mush into scripts:

```
from argparse import ArgumentParser, Namespace
from .configparser import RawConfigParser
from mush import Runner, requires, attr, item, returns
import logging, os, sqlite3, sys

log = logging.getLogger()
```

```

@requires (ArgumentParser)
def base_options(parser):
    parser.add_argument('config', help='Path to .ini file')
    parser.add_argument('--quiet', action='store_true',
                        help='Log less to the console')
    parser.add_argument('--verbose', action='store_true',
                        help='Log more to the console')

@requires (ArgumentParser)
def parse_args(parser):
    return parser.parse_args()

@requires (Namespace)
@returns ('config')
def parse_config(args):
    config = RawConfigParser()
    config.read(args.config)
    return dict(config.items('main'))

def setup_logging(log_path, quiet=False, verbose=False):
    handler = logging.FileHandler(log_path)
    handler.setLevel(logging.DEBUG)
    log.addHandler(handler)
    if not quiet:
        handler = logging.StreamHandler(sys.stderr)
        handler.setLevel(logging.DEBUG if verbose else logging.INFO)
        log.addHandler(handler)

class DatabaseHandler:
    def __init__(self, db_path):
        self.conn = sqlite3.connect(db_path)
    def __enter__(self):
        return self
    def __exit__(self, type, obj, tb):
        if type:
            log.exception('Something went wrong')
            self.conn.rollback()

```

We start with a function that adds the options needed by all our scripts to an `argparse` parser. This uses the `requires` decorator to tell Mush that it must be called with an `ArgumentParser` instance. See [Configuring Resources](#) for more details.

Next, we have a function that calls `parse_args()` on the parser and returns the resulting `Namespace`. The `parse_config()` function reads configuration from a file specified as a command line argument and so requires the `Namespace`. Since the config is a `dict`, we configure it as a named rather than typed resource. See [Named resources](#) for more details.

Finally, there is a function that configures log handling and a context manager that provides a database connection and handles exceptions that occur by logging them and aborting the transaction. Context managers like this are handled by Mush in a specific way, see [Context manager resources](#) for more details.

Each of these components can be separately and thoroughly tested. The Mush decorations are inert to all but the Mush `Runner`, meaning that they can be used independently of Mush in whatever way is convenient. As an example, the following tests use a `TempDirectory` and a `LogCapture` to fully test the database-handling context manager:

```

class DatabaseHandlerTests (TestCase):

    def setUp(self):

```

```
self.dir = TempDirectory()
self.addCleanup(self.dir.cleanup)
self.db_path = self.dir.getpath('test.db')
self.conn = sqlite3.connect(self.db_path)
self.conn.execute('create table notes '
                  '(filename varchar, text varchar)')
self.conn.commit()
self.log = LogCapture()
self.addCleanup(self.log.uninstall)

def test_normal(self):
    with DatabaseHandler(self.db_path) as handler:
        handler.conn.execute('insert into notes values (?, ?)',
                              ('test.txt', 'a note'))
        handler.conn.commit()
        # check the row was inserted and committed
        curs = self.conn.cursor()
        curs.execute('select * from notes')
        self.assertEqual(curs.fetchall(), [('test.txt', 'a note')])
        # check there was no logging
        self.log.check()

def test_exception(self):
    with ShouldRaise(Exception('foo')):
        with DatabaseHandler(self.db_path) as handler:
            handler.conn.execute('insert into notes values (?, ?)',
                                  ('test.txt', 'a note'))
            raise Exception('foo')
        # check the row not inserted and the transaction was rolled back
        curs = handler.conn.cursor()
        curs.execute('select * from notes')
        self.assertEqual(curs.fetchall(), [])
        # check the error was logged
        self.log.check(('root', 'ERROR', 'Something went wrong'))
```

2.2 Writing the specific parts of your script

Now that all the re-usable parts of the script have been abstracted, writing a specific script becomes a case of writing just two functions:

```
def args(parser):
    parser.add_argument('path', help='Path to the file to process')

def do(conn, path):
    filename = os.path.basename(path)
    with open(path) as source:
        conn.execute('insert into notes values (?, ?)',
                     (filename, source.read()))
```

As you can imagine, this much smaller set of code is simpler to test and easier to maintain.

2.3 Assembling the components into a script

So, we now have a library of re-usable components and the specific callables we require for the current script. All we need to do now is assemble these parts into the final script. The full details of this are covered in [Constructing runners](#) but two common patterns are covered below.

2.3.1 Cloning

With this pattern, a “base runner” is created, usually in the same place that other re-usable parts of the original script are located:

```
from mush import Runner, requires, attr, item, returns
base_runner = Runner(ArgumentParser)
base_runner.add(base_options, label='args')
base_runner.extend(parse_args, parse_config)
base_runner.add(setup_logging, requires(
    log_path = item('config', 'log'),
    quiet = attr(Namespace, 'quiet'),
    verbose = attr(Namespace, 'verbose')
))
```

The code above shows how to label a point, *args* in this case, enabling callables to be inserted at that point at a later time. See [Labels](#) for full details. It also shows some different ways of getting Mush to pass parts of an object returned from a previous callable to the parameters of another callable. See [Using parts of a resource](#) for full details.

Now, for each specific script, the base runner is cloned and the script-specific parts added to the clone leaving a callable that can be put in the usual block at the bottom of the script:

```
main = base_runner.clone()
main['args'].add(args, requires=ArgumentParser)
main.add(DatabaseHandler, requires=item('config', 'db'))
main.add(do,
    requires(attr(DatabaseHandler, 'conn'), attr(Namespace, 'path')))

if __name__ == '__main__':
    main()
```

2.3.2 Using a factory

This pattern is most useful when you have lots of scripts that all follow a similar pattern when it comes to assembling the runner from the common parts and the specific parts. For example, if all your scripts take a path to a config file and a path to a file that needs processing, you can write a factory function that returns a runner based on the callable that does the work as follows:

```
def options(parser):
    parser.add_argument('config', help='Path to .ini file')
    parser.add_argument('--quiet', action='store_true',
        help='Log less to the console')
    parser.add_argument('--verbose', action='store_true',
        help='Log more to the console')
    parser.add_argument('path', help='Path to the file to process')

def make_runner(do):
    runner = Runner(ArgumentParser)
    runner.add(options, requires=ArgumentParser)
```

```
runner.add(parse_args, requires=ArgumentParser)
runner.add(parse_config, requires=Namespace)
runner.add(setup_logging, requires(
    log_path = item('config', 'log'),
    quiet = attr(Namespace, 'quiet'),
    verbose = attr(Namespace, 'verbose')
))
runner.add(DatabaseHandler, requires=item('config', 'db'))
runner.add(
    do,
    requires(attr(DatabaseHandler, 'conn'), attr(Namespace, 'path'))
)
return runner
```

With this in place, the specific script becomes the `do()` function we abstracted above and a very short call to the factory:

```
main = make_runner(do)
```

A combination of the clone and factory patterns can also be used to get the best of both worlds. Setting up several base runners that are clones of a parent runner and having factories that take common callable patterns and return complete runners can be very powerful.

2.4 Testing

The examples above have shown how using Mush can make it easier to have smaller components that are easier to re-use and test, however care should still be taken with testing. In particular, it's a good idea to have some intergration tests that excercise the whole runner checking that it behaves as expected when all command line options are specified and when just the defaults are used.

When using the factory pattern, the factories themselves should be unit tested. It can also make tests easier to write by having a “testing runner” that sets up the required resources, such as database connections, while maybe doing some things differently such as not reading a configuration file from disk or using a `LogCapture` instead of file or stream log handlers.

Some specific tools that Mush provides to aid automated testing are covered in [Testing](#).

API Reference

class `mush.Runner` (**objects*)

A chain of callables along with declarations of their required and returned resources along with tools to manage the order in which they will be called.

`__add__` (*other*)

Return a new *Runner* containing the contents of the two *Runner* instances being added together.

`__call__` (*context=None*)

Execute the callables in this runner in the required order storing objects that are returned and providing them as arguments or keyword parameters when required.

A runner may be called multiple times. Each time a new *Context* will be created meaning that no required objects are kept between calls and all callables will be called each time.

Parameters `context` – Used for passing a context when context managers are used. You should never need to pass this parameter.

`__getitem__` (*label*)

Retrieve a *Modifier* for a previous labelled point in the runner.

add (*obj, requires=None, returns=None, label=None*)

Add a callable to the runner.

Parameters

- **obj** – The callable to be added.
- **requires** – The resources to required as parameters when calling *obj*. These can be specified by passing a single type, a string name or a *requires* object.
- **returns** – The resources that *obj* will return. These can be specified as a single type, a string name or a *returns*, *returns_mapping*, *returns_sequence* object.
- **label** – If specified, this is a string that adds a label to the point where *obj* is added that can later be retrieved with *Runner.__getitem__* ().

clone (*start_label=None, end_label=None, include_start=False, include_end=False*)

Return a copy of this *Runner*.

Parameters

- **start_label** – An optional string specifying the point at which to start cloning.
- **end_label** – An optional string specifying the point at which to stop cloning.
- **include_start** – If `True`, the point specified in *start_label* will be included in the cloned runner.

- **include_end** – If `True`, the point specified in `end_label` will be included in the cloned runner.

extend (**objs*)

Add the specified callables to this runner.

If any of the objects passed is a [Runner](#), the contents of that runner will be added to this runner.

replace (*original, replacement*)

Replace all instances of one callable with another.

No changes in requirements or call ordering will be made.

class `mush.requires` (**args, **kw*)

Represents requirements for a particular callable.

The passed in *args* and *kw* should map to the types, including any required [how](#), for the matching arguments or keyword parameters the callable requires.

String names for resources must be used instead of types where the callable returning those resources is configured to return the named resource.

__iter__ ()

When iterated over, yields tuples representing individual types required by arguments or keyword parameters in the form (`keyword_name`, `decorated_type`).

If the keyword name is `None`, then the type is for a positional argument.

class `mush.optional` (*type, *names*)

A [how](#) that indicates the callable requires the wrapped requirement only if it's present in the [Context](#).

class `mush.returns_result_type`

Default declaration that indicates a callable's return value should be used as a resource based on the type of the object returned.

`None` is ignored as a return value.

class `mush.returns_mapping`

Declaration that indicates a callable returns a mapping of type or name to resource.

class `mush.returns_sequence`

Declaration that indicates a callable's returns a sequence of values that should be used as a resources based on the type of the object returned.

Any `None` values in the sequence are ignored.

class `mush.returns` (**args*)

Declaration that specified names for returned resources or overrides the type of a returned resource.

This declaration can be used to indicate the type or name of a single returned resource or, if multiple arguments are passed, that the callable will return a sequence of values where each one should be named or have its type overridden.

class `mush.attr` (*type, *names*)

A [how](#) that indicates the callable requires the named attribute from the decorated type.

class `mush.item` (*type, *names*)

A [how](#) that indicates the callable requires the named item from the decorated type.

class `mush.context.Context`

Stores resources for a particular run.

add (*it, type*)

Add a resource to the context.

Optionally specify the type to use for the object rather than the type of the object itself.

exception `mush.context.ContextError` (*text, point=None, context=None*)

Errors likely caused by incorrect building of a runner.

class `mush.modifier.Modifier` (*runner, callpoint, label=None*)

Used to make changes at a particular point in a runner. These are returned by `Runner.add()` and `Runner.__getitem__()`.

add (*obj, requires=None, returns=None, label=None*)

Parameters

- **obj** – The callable to be added.
- **requires** – The resources to required as parameters when calling *obj*. These can be specified by passing a single type, a string name or a *requires* object.
- **returns** – The resources that *obj* will return. These can be specified as a single type, a string name or a *returns*, *returns_mapping*, *returns_sequence* object.
- **label** – If specified, this is a string that adds a label to the point where *obj* is added that can later be retrieved with `Runner.__getitem__()`.

If no label is specified but the point which this *Modifier* represents has any labels, those labels will be moved to the newly inserted point.

add_label (*label, callpoint=None*)

Add a label to the point represented by this *Modifier*.

Parameters **callpoint** – For internal use only.

class `mush.declarations.how` (*type, *names*)

The base class for type decorators that indicate which part of a resource is required by a particular callable.

Parameters

- **type** – The resource type to be decorated.
- **names** – Used to identify the part of the resource to extract.

process (*o*)

Extract the required part of the object passed in. *missing* should be returned if the required part cannot be extracted. *missing* may be passed in and is usually be handled by returning *missing* immediately.

`mush.declarations.nothing = requires()`

A singleton *requires* indicating that a callable has no required arguments.

`mush.declarations.result_type = returns_result_type()`

A singleton indicating that a callable's return value should be stored based on the type of that return value

For details of how to install the package or get involved in its development, please see the sections below:

Installation Instructions

The easiest way to install Mush is:

```
pip install mush
```

Python version requirements

This package has been tested with Python 2.7, 3.3+ on Linux, but is also expected to work on Mac OS X and Windows.

Development

This package is developed using continuous integration which can be found here:

<https://travis-ci.org/Simplistix/mush>

The latest development version of the documentation can be found here:

<http://mush.readthedocs.org/en/latest/>

If you wish to contribute to this project, then you should fork the repository found here:

<https://github.com/Simplistix/mush/>

Once that has been done and you have a checkout, you can follow these instructions to perform various development tasks:

5.1 Setting up a virtualenv

The recommended way to set up a development environment is to turn your checkout into a virtualenv and then install the package in editable form as follows:

```
$ virtualenv .  
$ bin/pip install -U -e .[test,build]
```

5.2 Running the tests

Once you have a buildout, the tests can be run as follows:

```
$ bin/nosetests
```

5.3 Building the documentation

The Sphinx documentation is built by doing the following from the directory containing setup.py:

```
$ cd docs  
$ make html
```

5.4 Making a release

To make a release, just update the version in `setup.py`, update the change log, tag it and push to <https://github.com/Simplistix/mush> and Travis CI should take care of the rest.

Once Travis CI is done, make sure to go to <https://readthedocs.org/projects/mush/versions/> and make sure the new release is marked as an Active Version.

Changes

6.1 2.1 (14 December 2015)

- Typo fixes in documentation.
- Indicate that Python 2.6 is no longer supported.
- Raise exceptions when arguments to `requires()` and `returns()` are not either types or labels.
- Allow tuples and lists to be passed to `add()`, they will automatically be turned into a `requires()` or `returns()`.
- Better error messages when a requirement is not found in the `Context`.

Thanks to Dani Fortunov for the documentation review.

6.2 2.0 (11 December 2015)

- Re-write dropping all the heuristic callable ordering in favour of building up defined sequences of callables with labelled insertion points.

6.3 1.3 (21 October 2015)

- Official support for Python 3.
- Drop official support for Windows, although things should still work.
- Move to Travis CI, Read The Docs and Coveralls for development.
- ‘How’ decorators like `attr()` and `item()` can now be nested as well as individually performing nested actions.
- Add `returns()` and `add_returning()` as new ways to override the type of a returned value.
- A better pattern for “marker types”.

6.4 1.2 (11 December 2013)

- Use `nothing` instead of `None` for marker return types, fixing a bug that occurred when a callable tried to type-map a result that was `None`.

- Add an `after()` type wrapper for callables that need to wait until after a resource is used but that can't accept that resource as a parameter.

6.5 1.1 (27 November 2013)

- Allow runners to be instantiated using other runners.
- Allow `Runner.extend()` to be passed `Runner` instances.
- Allow `requires()` decorations to be stacked.
- Add a `Runner.replace()` method to aid with testing assembled runners.

6.6 1.0 (29 October 2013)

- Initial Release

License

Copyright (c) 2013 Simplistix Ltd, 2015 Chris Withers

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the “Software”), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Indices and tables

- `genindex`
- `modindex`
- `search`

m

`mush`, [21](#)

`mush.context`, [22](#)

`mush.declarations`, [23](#)

`mush.modifier`, [23](#)

Symbols

`__add__()` (mush.Runner method), 21
`__call__()` (mush.Runner method), 21
`__getitem__()` (mush.Runner method), 21
`__iter__()` (mush.requires method), 22

A

`add()` (mush.context.Context method), 22
`add()` (mush.modifier.Modifier method), 23
`add()` (mush.Runner method), 21
`add_label()` (mush.modifier.Modifier method), 23
`attr` (class in mush), 22

C

`clone()` (mush.Runner method), 21
`Context` (class in mush.context), 22
`ContextError`, 23

E

`extend()` (mush.Runner method), 22

H

`how` (class in mush.declarations), 23

I

`item` (class in mush), 22

M

`Modifier` (class in mush.modifier), 23
`mush` (module), 21
`mush.context` (module), 22
`mush.declarations` (module), 23
`mush.modifier` (module), 23

N

`nothing` (in module mush.declarations), 23

O

`optional` (class in mush), 22

P

`process()` (mush.declarations.how method), 23

R

`replace()` (mush.Runner method), 22
`requires` (class in mush), 22
`result_type` (in module mush.declarations), 23
`returns` (class in mush), 22
`returns_mapping` (class in mush), 22
`returns_result_type` (class in mush), 22
`returns_sequence` (class in mush), 22
`Runner` (class in mush), 21