

---

# **Chopsticks Documentation**

*Release 0.5*

**Daniel Pope**

**Jun 07, 2017**



---

## Contents

---

<b>1</b>	<b>Introduction</b>	<b>3</b>
<b>2</b>	<b>Python 2/3</b>	<b>5</b>
<b>3</b>	<b>Jupyter Notebooks</b>	<b>7</b>
<b>4</b>	<b>How it works</b>	<b>9</b>
<b>5</b>	<b>Chopsticks vs ...</b>	<b>11</b>
5.1	Ansible . . . . .	11
5.2	Fabric . . . . .	11
<b>6</b>	<b>Tunnels</b>	<b>13</b>
6.1	Tunnel reference . . . . .	13
6.2	Writing new tunnels . . . . .	14
<b>7</b>	<b>Groups</b>	<b>17</b>
7.1	Group API . . . . .	17
7.2	Results . . . . .	18
7.3	Examples . . . . .	18
<b>8</b>	<b>How to...</b>	<b>21</b>
8.1	How to write a single-file Chopsticks script . . . . .	21
8.2	How to customise interpreter paths . . . . .	22
<b>9</b>	<b>Indices and tables</b>	<b>23</b>



Chopsticks is an orchestration and remote execution library. It lets you run Python code elsewhere: on remote hosts over SSH, in a Docker sandbox, on the local host (optionally with `sudo`) - even all of these in parallel.

Naturally this is agentless and nothing needs to be installed on remote hosts except Python and an SSH agent.

Chopsticks was built for extensibility. Remote hosts may import Python code from the orchestration host on demand, so remote agents can immediately use new functions you define. In effect, you have access to the same codebase on remote hosts as on the orchestration host.

Contents:



# CHAPTER 1

---

## Introduction

---

With chopsticks you can simply import functions and hand them to the remote host to be executed.

First stand up an SSH Tunnel:

```
from chopsticks.tunnel import Tunnel
tun = Tunnel('troy.example.com')
```

Then you can pass a function, to be called on the remote host:

```
import time
print('Time on %s:' % tun.host, tun.call(time.time))
```

You can use any pure-Python function in the current codebase, meaning you can create your own libraries of orchestration functions to call on remote hosts (as well as functions that call out to remote hosts using Chopsticks). Naturally those functions can import pure-Python libraries and so on. Your entire local codebase should just work remotely.

Group allows for executing a callable on a number of hosts in parallel:

```
from chopsticks.group import Group

group = Group([
    'web1.example.com',
    'web2.example.com',
    'web3.example.com',
])
for host, t in group.call(time.time).successful():
    print('Time on %s:' % host, t)
```

You can also run your code within Docker containers:

```
from chopsticks.tunnel import Docker
from chopsticks.facts import python_version

dkr = Docker('py36', image='python:3.6')
print(dkr.call(python_version))
```

Tunnels and Groups connect lazily (or you can connect them proactively by calling `connect()`). They are also usable as context managers:

```
# Explicitly connect and disconnect
group.connect()
group.call(time.time)
group.close()

# Reconnect and disconnect as context manager
with group:
    group.call(time.time)

# Implicit reconnect
group.call(time.time)

# Disconnect when destroyed
del group
```

Naturally, any remote state (imports, globals, etc) is lost when the Tunnel/Group is closed.



## CHAPTER 2

---

### Python 2/3

---

Chopsticks supports both Python 2 and Python 3.

Because Chopsticks takes the view that agents run out of the same codebase as the controller, agents will attempt to use a similar Python interpreter to the one for the controller process:

- `/usr/bin/python2` if the controller process is (any) Python 2.
- `/usr/bin/python3` if the controller process is (any) Python 3.



## CHAPTER 3

---

### Jupyter Notebooks

---

For interactive exploration, Chopsticks can also be used within [Jupyter Notebooks](#). Functions defined in Notebook cells are sent over the tunnel as fragments of Python source (rather than imported).

This generally gives good results, but is somewhat more magical than Chopsticks' standard import behaviour. Any odd behaviour should be reported via the [issue tracker](#).



## CHAPTER 4

---

### How it works

---

The SSH tunnel invokes the `python` binary on the remote host, and feeds it a bootstrap script via `stdin`.

Once bootstrapped, the remote “agent” sets up bi-directional communication over the `stdin/stdout` of the tunnel. This communication is used (currently) for two purposes:

- An RPC system to invoke arbitrary callables within the remote agent and pass the returned values back to the controller.
- A PEP-302 import hook system, allowing the remote agent to import pure-Python code from the controller (NB. the controller can only serve Python modules that live within the filesystem - import hooks such as `zipimport/compressed eggs` are not currently supported).

`stdin/stdout` on the agent are redirected to `/dev/null`, so calling `print()` on the remote machine will not break the tunnel.

`stderr` is echoed to the controlling console, prefixed with a hostname to identify which Tunnel it issued from. This can therefore be used to feed debugging information back to the orchestration host.



---

## Chopsticks vs ...

---

It's natural to draw comparisons between Chopsticks and various existing tools, but Chopsticks is a library, not an orchestration framework in its own right, and other tools could potentially build on it.

### Ansible

Ansible's YAML syntax is a lot more restrictive than Python. It is friendly for simple cases, but becomes increasingly ugly and convoluted as your scripts become more complex. By writing your orchestration scripts in Python you can take advantage of Python's rich ecosystem of syntax and tools for writing clean Python code and documenting it, which apply even for very complicated use cases.

Ansible's remote execution model involves dropping scripts, calling them, and deleting them. In Ansible 2.1, some of Ansible's support code for Python-based Ansible plugins gets shipped over SSH as part of a zipped bundle; but this doesn't extend to your own code extensions. So Chopsticks is more easily and naturally extensible: write your code how you like and let Chopsticks deal with getting it running on the remote machine.

### Fabric

The big difference between [Fabric](#) and Chopsticks is that Fabric will only execute shell commands on the remote host, not Python callables. Of course you can drop Python scripts and call them, but then you're back in Ansible territory for extensibility, or you have to bootstrap the dependencies needed to execute such scripts manually.

The difference in concept goes deeper: Fabric tries to be "of SSH", exploiting all the cool SSH tunnelling features. Chopsticks doesn't care about SSH specifically; it only cares about Python and pipes. This is what allows it to work identically with Docker or subprocesses as with remote SSH hosts.





Tunnels are the lowest-level API, used for invoking commands on an individual host or container. For a higher-level API that allows invoking commands in parallel across a range of hosts, see *Groups*.

An established tunnel can be used to invoke commands and receive results.

## Tunnel reference

### SSH

**class** `chopsticks.tunnel.SSHTunnel` (*host*, *user=None*, *sudo=False*)

A tunnel that connects to a remote host over SSH.

**call** (*callable*, *\*args*, *\*\*kwargs*)

Call the given callable on the remote host.

The callable must return a value that can be serialised as JSON, but there is no such restriction on the parameters.

**fetch** (*remote\_path*, *local\_path=None*)

Fetch one file from the remote host.

If *local\_path* is given, it is the local path to write to. Otherwise, a temporary filename will be used.

This operation supports arbitrarily large files (file data is streamed, not buffered in memory).

The return value is a dict containing:

- *local\_path* - the local path written to
- *remote\_path* - the absolute remote path
- *size* - the number of bytes received
- *shalsum* - a sha1 checksum of the file data

**put** (*local\_path*, *remote\_path=None*, *mode=420*)

Copy a file to the remote host.

If *remote\_path* is given, it is the remote path to write to. Otherwise, a temporary filename will be used.

*mode* gives is the permission bits of the file to create, or 0o644 if unspecified.

This operation supports arbitrarily large files (file data is streamed, not buffered in memory).

The return value is a dict containing:

- *remote\_path* - the absolute remote path
- *size* - the number of bytes received
- *shalsum* - a sha1 checksum of the file data

`chopsticks.tunnel.Tunnel`  
alias of `SSHTunnel`

## Docker

**class** `chopsticks.tunnel.Docker` (*name*, *image='python:2.7'*, *rm=True*)

A tunnel connected to a throwaway Docker container.

## Subprocess

**class** `chopsticks.tunnel.Local` (*name='localhost'*)

A tunnel to a subprocess on the same host.

## Sudo

**class** `chopsticks.tunnel.Sudo` (*user='root'*, *name=None*)

A tunnel to a process on the same host, launched with sudo.

The *Sudo* tunnel does not deal with password dialogues etc. In order for this to work you must configure `sudo` not to need a password. You can do this with these lines in `/etc/sudoers`:

```
Cmdnd_Alias PYTHON_CMDS = /usr/bin/python, /usr/bin/python2, /usr/bin/python3
%somegroup ALL=NOPASSWD: PYTHON_CMDS
```

This would allow users in the group `somegroup` to be able to run the system Python interpreters using `sudo`, without passwords.

**Warning:** Naturally, as Chopsticks is a framework for executing arbitrary code, this allows executing arbitrary code as root. Only make this change if you are happy with relaxing security in this way.

## Writing new tunnels

It is possible to write a new tunnel driver for any system that allows you to execute a `python` binary with direct relay of `stdin` and `stdout` pipes. To do this, simply subclass `chopsticks.group.SubprocessTunnel`. Note that all tunnel instances must have a `host` attribute which is used as the key for the result in the `GroupResult` dictionary when executing tasks in a `Group`.

So, strictly, these requirements apply:

- The tunnel setup machinery should not write to `stdout` - else you will have to identify and consume this output.
- The tunnel setup machinery should not read from `stdin` - else you will have to feed the required input.
- Both `stdin` and `stdout` must be binary-safe pipes.

The tunnel machinery may write to `stderr`; this output will be presented to the user.



Groups can be used to perform a remote operation in parallel across a number of hosts, and collect the results.

## Group API

**class** `chopsticks.group.Group` (*hosts*)

A group of hosts, for performing operations in parallel.

**\_\_init\_\_** (*hosts*)

Construct a group from a list of tunnels or hosts.

*hosts* may contain hostnames - in which case the connections will be made via SSH using the default settings. Alternatively, it may contain tunnel instances.

**call** (*callable*, \**args*, \*\**kwargs*)

Call the given callable on all hosts in the group.

The given callable and parameters must be pickleable.

However, the callable's return value has a tighter restriction: it must be serialisable as JSON, in order to ensure the orchestration host cannot be compromised through pickle attacks.

The return value is a *GroupResult*.

**fetch** (*remote\_path*, *local\_path=None*)

Fetch files from all remote hosts.

If *local\_path* is given, it is a local path template, into which the tunnel's `host` name will be substituted using `str.format()`. Hostnames generated in this way must be unique.

For example:

```
group.fetch('/etc/passwd', local_path='passwd-{host}')
```

If *local\_path* is not given, a temporary file will be used for each host.

Returns a *GroupResult* of dicts, each containing:

- `local_path` - the local path written to
- `remote_path` - the absolute remote path
- `size` - the number of bytes received
- `shalsum` - a sha1 checksum of the file data

**put** (*local\_path*, *remote\_path=None*, *mode=420*)  
Copy a file to all remote hosts.

If `remote_path` is given, it is the remote path to write to. Otherwise, a temporary filename will be used (which will be different on each host).

`mode` gives the permission bits of the files to create, or 0o644 if unspecified.

This operation supports arbitrarily large files (file data is streamed, not buffered in memory).

The return value: class: *GroupResult* of dicts, each containing:

- `remote_path` - the absolute remote path
- `size` - the number of bytes received
- `shalsum` - a sha1 checksum of the file data

## Results

**class** `chopsticks.group.GroupResult`

The results of a `Group.call()` operation.

`GroupResult` behaves as a dictionary of results, keyed by hostname, although failures from individual hosts are represented as `ErrorResult` objects.

Methods are provided to easily process successes and failures separately.

**failures** ()

Iterate over failed results as (host err) pairs.

**successful** ()

Iterate over successful results as (host, value) pairs.

**class** `chopsticks.group.ErrorResult` (*msg*, *tb=None*)

Indicates an error returned by the remote host.

Because tracebacks or error types cannot be represented across hosts this will simply consist of a message.

Error results provide the following attributes:

**msg**

A human-readable error message.

**tb**

The traceback from the remote host as a string, or `None` if unavailable.

## Examples

For example, this code:

```
from chopsticks.facts import ip
from chopsticks.group import Group

group = Group([
    'web1.example.com',
    'web2.example.com',
    'web3.example.com',
])
for host, addr in group.call(ip).items():
    print('%s ip:' % host, addr)
```

might output:

```
web1.example.com ip: 196.168.10.5
web3.example.com ip: 196.168.10.7
web2.example.com ip: 196.168.10.6
```

You could also construct a group from existing tunnels - or mix and match:

```
all_hosts = Group([
    'web1.example.com',
    Docker('example'),
    Local('worker')
])
```





## How to write a single-file Chopsticks script

Chopsticks will work very well with a neatly organised codebase of management functions, but you can also write a single file script.

Chopsticks has special logic to handle this case, which is different from the standard import machinery.

The cleanest way to write this script would be:

```
from chopsticks import Tunnel

def do_it():
    return 'done'

if __name__ == '__main__':
    with Tunnel('remote') as tun:
        tun.call(do_it)
```

Actually, only the `do_it()` function, and various globals it uses, are sent to the remote host. This code will work just fine:

```
from chopsticks import Tunnel

def do_it():
    return 'done'

tunnel = Tunnel('remote')
tunnel.call(do_it)
```

This also allows Chopsticks to be used from within [Jupyter Notebooks](#).

## How to customise interpreter paths

Chopsticks assumes that the interpreter path on a remote host will be `/usr/share/python2` for Python 2 and `/usr/share/python3` for Python 3. However, these paths may not always be correct.

To override the path of the interpreter you can simple subclass `Tunnel` (or the tunnel type you wish to use), and modify the `python2` and `python3` class attributes:

```
class MyTunnel(Tunnel):  
    python3 = '/usr/local/bin/python2'
```

To do this for all tunnels of the same type, modify the attribute on the type:

```
Tunnel.python2 = '/usr/bin/python2'
```

## CHAPTER 9

---

### Indices and tables

---

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



## Symbols

`__init__()` (chopsticks.group.Group method), 17

### C

`call()` (chopsticks.group.Group method), 17

`call()` (chopsticks.tunnel.SSHTunnel method), 13

### D

Docker (class in chopsticks.tunnel), 14

### E

ErrorResult (class in chopsticks.group), 18

### F

`failures()` (chopsticks.group.GroupResult method), 18

`fetch()` (chopsticks.group.Group method), 17

`fetch()` (chopsticks.tunnel.SSHTunnel method), 13

### G

Group (class in chopsticks.group), 17

GroupResult (class in chopsticks.group), 18

### L

Local (class in chopsticks.tunnel), 14

### M

`msg` (chopsticks.group.ErrorResult attribute), 18

### P

`put()` (chopsticks.group.Group method), 18

`put()` (chopsticks.tunnel.SSHTunnel method), 13

### S

SSHTunnel (class in chopsticks.tunnel), 13

`successful()` (chopsticks.group.GroupResult method), 18

Sudo (class in chopsticks.tunnel), 14

### T

`tb` (chopsticks.group.ErrorResult attribute), 18

Tunnel (in module chopsticks.tunnel), 14