
plumbery Documentation

Release 1.0.0

Bernard Paques

July 15, 2016

1	Plumbery	1
1.1	Features	1
1.2	Contributors	2
1.3	Credits	2
2	Why plumbery?	3
2.1	One day in the life of an ordinary infrastructure manager	3
2.2	Infrastructure as code	4
2.3	Infrastructure as code at Dimension Data with Apache Libcloud	4
2.4	What should be described in the fittings plan?	5
2.5	How to run Plumbery?	5
3	Using plumbery	9
4	How to install plumbery?	11
4.1	Installing plumbery package	11
4.2	Installing plumbery from GitHub	12
4.3	Configure a Linux machine, or a Mac	13
4.4	Configure a Windows machine	13
4.5	Test the installation	13
4.6	Do a quick deployment	13
5	The structure and content of fittings file	15
5.1	Keywords used by plumbery	15
5.2	Multiple documents in one fittings file	15
5.3	Deploying in multiple geographies	16
5.4	Combining private and public clouds in a deployment	16
5.5	Complete example	16
6	How to contribute to Plumbery?	19
6.1	You are not a developer? We are glad that you are involved	19
6.2	Ready to contribute? Here's how to set up Plumbery for local development	20
6.3	Pull Request Guidelines	21
7	History	23
7.1	1.0.0 (2016-07-15)	23
7.2	0.9.9 (2016-06-10)	23
7.3	0.9.8 (2016-06-10)	23
7.4	0.9.7 (2016-06-08)	23

7.5	0.9.6 (2016-06-06)	23
7.6	0.9.5 (2016-05-28)	24
7.7	0.9.4 (2016-05-27)	24
7.8	0.9.3 (2016-05-27)	24
7.9	0.9.2 (2016-05-27)	24
7.10	0.9.1 (2016-05-26)	24
7.11	0.9.0 (2016-05-25)	24
7.12	0.8.1 (2016-05-23)	24
7.13	0.8.0 (2016-04-13)	25
7.14	0.7.0 (2016-04-06)	25
7.15	0.6.0 (2016-03-15)	25
7.16	0.5.0 (2016-03-13)	25
7.17	0.4.3 (2016-02-28)	26
7.18	0.4.2 (2016-02-14)	26
7.19	0.4.1 (2016-01-24)	27
7.20	0.4.0 (2016-01-17)	27
7.21	0.3.4 (2016-01-06)	28
7.22	0.3.3 (2015-12-31)	28
7.23	0.3.2 (2015-12-23)	28
7.24	0.2.4 (2015-12-13)	29
7.25	0.2.3 (2015-12-07)	29
7.26	0.2.2 (2015-12-05)	30
7.27	0.2.1 (2015-12-02)	30
7.28	0.2.0 (2015-11-29)	30
7.29	0.1.2 (2015-11-27)	30
7.30	0.1.0 (2015-11-20)	31
8	plumbery package	33
8.1	Subpackages	33
8.2	Submodules	33
8.3	Module contents	34
	Python Module Index	35

Plumber

Infrastructure as code at Dimension Data with Apache Libcloud

- Documentation: [Plumber at ReadTheDocs](#)
- Python package: [Plumber at PyPi](#)
- Source code: [Plumber at GitHub](#)
- Free software: [Apache License \(2.0\)](#)

1.1 Features

- 41 tutorials covering popular use cases
- Read fittings plan in YAML
- Load parameters in separate YAMI file
- Use cloud API to create the network infrastructure, and to build nodes
- Bootstrap nodes with cloud-init configuration directives
- Inject actual addresses and on-the-fly secrets to contextualisation
- Create RSA keys and passwords to secure deployments
- Preserve random secrets across multiple invocations
- Create network domains and Ethernet networks
- Reserve public IPv4 addresses
- Manage network address translation rules
- Manage firewall rules
- Create load balancers with pools of nodes
- All images in libraries are available to new nodes
- Specify number of CPU, or core per CPU, and CPU speed
- Specify node memory
- Add virtual disks and specify tiers of storage
- Add multiple network interfaces to nodes
- Set private IPv4 statically

- Add public IPv4 addresses to nodes
- Add monitoring to nodes
- Add backup to nodes
- Build all blueprints
- Build a blueprint across multiple locations
- Start all nodes
- Start nodes belonging to the same blueprint
- Polish Linux nodes for quick bootstrapping
- Build a full inventory of nodes that have been deployed
- Reflect fittings into a ready-to-use inventory for ansible
- Stop all nodes
- Stop nodes belonging to the same blueprint
- Wipe all nodes
- Wipe nodes belonging to the same blueprint
- Destroy part of the setup, or all blueprints
- You can extend plumbery with your own software, it has been designed for that
- Run from the command line, or as a python library
- Accept custom actions from the command line too
- Limit plumbing to some locations if needed

1.2 Contributors

- [Bernard Paques](#) [Dimension Data employee] (development lead)
- [Anthony Shaw](#) [Dimension Data employee]
- [Olivier Grosjeanne](#) [Dimension Data employee]
- [Jacques Clément](#) [Dimension Data employee]
- [Asim Khawaja](#) [Dimension Data employee]

1.3 Credits


- Cloud-Init
- Apache Libcloud
- netifaces
- PyYAML
- Cookiecutter
- cookiecutter-pypackage

Why plumbery?

Plumbery is an open-source project that was initiated by Dimension Data to accelerate the digital transformation of our clients.

2.1 One day in the life of an ordinary infrastructure manager

This project started in response to a very common issue. How to accelerate the creation, the handling, and the destruction of up to 60 virtual servers easily? The purpose here is not to supplement Chef, Puppet, and other smart configuration management that already exist. The first priority is to industrialise operations on the underlying virtualised iron.



The Plumbery toolbox is based on a central description of servers, documented in a text file in **YAML**. If you do not know **YAML** yet, don't be afraid, this may be the most simple and user-friendly language for infrastructure managers.

Then very simple command is used to act on the infrastructure. For example you can build the entire set of resources described in the fittings plan directly from the command line:

```
$ python -m plumbery fittings.yaml deploy
```

This will load the YAML file, parse it, and call the cloud API to make it happen. Relax, and grab some coffee while plumbery adds network domains, Ethernet networks, and servers as per your specifications. The engine may plumb at various data centres spread on Earth, thanks to the power of Apache Libcloud. Nodes will be started and contextualised with cloud-init directives as well.

Now you can concentrate on important things, connect remotely to the nodes, play with them, run different tests, etc. At the end of the game, you would just have to stop all servers and destroy them as per following command:

```
$ python -m plumbery fittings.yaml dispose
```

2.2 Infrastructure as code

Plumbery fully supports the idea that infrastructure should be handled like code. For people who have relied for years on the power of physical stuff, this may be a shock. So let's repeat it again. When interesting resources have been virtualised and accessed remotely then suddenly infrastructure managers have new challenges. Infrastructure should be treated like code is. A lot of best practices are coming with this paradigm.

And issues, too. Becoming a trusted Ops –yes, the second half-part of DevOps– is not so easy. How to deal with developers while being not a software developer? Well, this is exactly why Plumbery has been developed.

Actually we believe that every infrastructure manager should have enough skills to manage sophisticated configurations with plumbery. Here we are referring to tasks like the following:

- open and edit textual files
- understand and modify configuration files
- connect to a Linux server via ssh
- set environment variables
- run a command from the prompt line
- execute a python program

2.3 Infrastructure as code at Dimension Data with Apache Libcloud

If infrastructure should be handled like code, then your first task is to document a target deployment into a text file. In the context of Plumbery, this is called the fittings plan, and it is commonly put into a file named `fittings.yaml`. This is also called the fittings file.

With that in hands, you can then ask Plumbery to act on the infrastructure and on nodes. The engine has built-in code to cover the full life cycle:

- deploy the infrastructure and configure it
- build nodes
- start nodes
- stop nodes
- destroy nodes
- dispose the infrastructure and release all resources

2.4 What should be described in the fittings plan?

Most, if not all, elements of Dimension Data cloud services, can be handled by plumbery. This is covering:

- Network Domains
- Networks
- Private IP addresses
- Public IP addresses
- NAT rules
- Firewall rules
- Cloud Servers
- Load-Balancers
- Monitoring dashboard
- Cloud Backup

Plumbery has been extended so that it can handle external cloud services and act within cloud nodes:

- terraform - orchestrate resources at AWS or ar Microsoft Azure
- cloud-init - configure the operating system, files, and applications

The fittings plan is written in YAML, and it must have multiple documents in it. The first document provides general configuration parameters for the engine. Subsequent documents describe the various locations for the fittings.

An example of a minimum fittings plan:

```
locationId: EU6
regionId: dd-eu

blueprints:

  - myBluePrint:
    domain:
      name: myDC
    ethernet:
      name: myVLAN
      subnet: 10.1.10.0
    nodes:
      - myServer
```

In this example, the plan is to deploy a single node in the data centre at Frankfurt, in Europe. The node *myServer* will be placed in a network named *myVLAN*, and the network will be part of a network domain acting as a virtual data centre, *myDC*. The blueprint has a name, *myBluePrint*, so that it can be handled independently from other blueprints.

To learn more: [The structure and content of fittings file](#)

2.5 How to run Plumbery?

2.5.1 Using plumbery from the command-line

Move first to the directory that contains your fittings plan, and then run:

```
$ python -m plumbery fittings.yaml deploy
```

Plumbery will load `fittings.yaml`, then build all blueprints there.

As you can expect, plumbery can be invoked through the entire life cycle of your fittings:

```
$ python -m plumbery fittings.yaml build
$ python -m plumbery fittings.yaml start
$ python -m plumbery fittings.yaml prepare

... nodes are up and running here ...

$ python -m plumbery fittings.yaml stop
$ python -m plumbery fittings.yaml destroy
```

The table below presents succinctly all actions that are supported by plumbery.

Action	Description
deploy	equivalent to: build + configure + start + prepare
dispose	equivalent to: stop + destroy
build	create network domains, networks, and nodes
configure	adds public IP addresses, NAT and firewall rules
start	start nodes
prepare	contextualise nodes via ssh and cloud-init
information	display information put in fittings plan
inventory	produce an inventory of all assets deployed
ansible	allow ansible to handle nodes and groups deployed by plumbery
ping	check the status of nodes and display network addresses
stop	stop nodes
wipe	destroy only nodes
destroy	destroy nodes and other resources
polish	apply all polishers configured in fittings plan
secrets	display secrets such as random passwords, etc.

How do I handle a subset of cloud resources?

You are right to ask. A common use case is the immediate creation of a new group of servers for some developer. You would not like this to interfere with existing servers. On the other hand, you may really want to destroy a group of unused servers that just add to the monthly invoices.

In plumbery the full fittings plan is split in multiple blueprints, and you can handle each of them separately. The most natural way to think about this is to conceive services as group of servers. For example, the blueprint `docker` is actually a cluster of nodes plugged into the same network. Look at the sample `fittings.yaml` file to get an idea of what you can put in a blueprint.

Then you can handle a single blueprint independently from the others:

```
$ python -m plumbery fittings.yaml build docker
$ python -m plumbery fittings.yaml build sql

... and then later ...

$ python -m plumbery fittings.yaml destroy sql
$ python -m plumbery fittings.yaml build mongodb
```

How to plumb only at a selected location?

By default Plumbery looks at every location mentioned in fittings plan. Sometimes you may want to limit actions performed to some locations. For this, mention the name of the target location, prefixed by @. As an example, here would be the command to build SQL servers only at NA12:

```
$ python -m plumbery fittings.yaml build sql @NA12
```

How to get help from the command line?

```
$ python -m plumbery -h
```

When you are not sure of directives put in fittings plan, you can ask plumbery for a dry-run. Use the -s switch to keep it safe:

```
$ python -m plumbery <your_various_args> -s
```

If some strange behaviour occurs, and you cannot understand what is happening, then you can use the debug option to get more information:

```
$ python -m plumbery <your_various_args> -d
```

Of course the switches -s and -d can be combined if needed.

2.5.2 Using from Docker

The plumbery Docker image is the simplest and easiest way to use Plumbery, you can process multiple deployments at the same time, isolate the processes and there is no requirement to install any additional software. to download the plumbery image simply run:

```
$ docker pull dimensiondataresearch/plumbery
```

This will download the latest image from <https://hub.docker.com/r/dimensiondataresearch/plumbery/>

You can then run the image with the following environment variables:

- **MCP_USERNAME** - your API username
- **MCP_PASSWORD** - your API password
- **SHARED_SECRET** - the password for the deployed servers
- **FITTINGS** - The URL to some fittings file, this can be a HTTP/HTTPS or a FTP/FTPS address. Look in <https://github.com/DimensionDataCBUSydney/plumbery-contrib> for examples

A fictitious example could be:

```
$ docker run -e "MCP_USERNAME=bob_smith" -e "MCP_PASSWORD=superPassword!" \
  -e "SHARED_SECRET=superPassword!" \
  -e "FITTINGS=https://raw.githubusercontent.com/DimensionDataCBUSydney/plumbery-contrib/master/fittings.yaml" \
  dimensiondataresearch/plumbery
```

If needed, you have more variables to play with:

- **WGET_OPTS** can be set for the fetching of the fittings file, the following options are supported: * “--no-check-certificate” to disable SSL certificate validation * “-ftp-user=user -ftp-password=password” for FTP credentials * “--no-passive-ftp” disable FTP passive transfer mode, for use with proxys * “-http-user=user -http-password=password” for HTTP basic authentication

- **OPTS** can be set for the plumbery specific options, i.e.: * “-d” to set debug mode and get more information * “-safe” to only run the fittings file but not make any changes, for testing
- **ACTION** is “deploy” by default, but can be set to any of the potential actions for plumbery

2.5.3 Using Plumbery as a python library

Since Plumbery is easy to load, you can use it interactively like in the following example:

```
>>>from plumbery.engine import PlumberyEngine
>>>PlumberyEngine('fittings.yaml').build_blueprint('beachhead control')
...

```

If you are writing some code using Plumbery as a library, you would import the engine and use it, as with any other python module. For example:

```
from plumbery.engine import PlumberyEngine

engine = PlumberyEngine('fittings.yaml')
engine.do('build', 'docker')
engine.do('start', 'docker')
engine.do('prepare', 'docker')
```

The source code is available on-line, check the [Plumbery repository at GitHub](#).

Using plumbing

How to install plumbery?

Adding plumbery as a new package on your workstation may be the simplest way to go. Consider the installation from git if you want to get all of the source code.

4.1 Installing plumbery package

You will need the latest version of the Python 2.7 language, which you can download for Linux, Mac and Windows at the following address:

```
http://python.org/download/releases/2.7/
```

You need pip, a package management system, to install modules related to plumbery. Should your system not have pip installed, run this command:

```
$ sudo apt-get install python-pip
```

4.1.1 Install Apache Libcloud

The orchestration of cloud services is a hot topic these days. Apache Libcloud is evolving swiftly thanks to many contributions. For this reason, you are encouraged to install the development version of the library:

```
$ sudo pip install -e git+https://git-wip-us.apache.org/repos/asf/libcloud.git@trunk#egg=apache-libcloud
```

This version is the one used by Plumbery, and it may be more advanced than the stable version of Apache Libcloud.

4.1.2 Install the plumbery package

Plumbery is available as a python package, so the installation, the upgrade, and the removal of the software are really easy.

Install the plumbery package

Plumbery is [available on PyPi](#). You can install latest stable version using pip:

```
$ sudo apt-get install python-pip python-dev
$ sudo pip install plumbery
```

The installation of *python-dev* is required for the installation of the module *netifaces*, that is used by Plumbery to get information about network interfaces.

For installation on Windows, you may need to first install the Python Compiler for VC++. <https://www.microsoft.com/en-us/download/confirmation.aspx?id=44266> Note this only works for Python 2.7. If you get an error on installation saying “error: Unable to find vcvarsall.bat” this indicates you need to install this package.

```
$ C:\Python27\Scripts\virtualenv.exe .
$ .\Script\pip install plumbery
```

Upgrade the plumbery package

Use the following command to retrieve the version of plumbery that has been installed on a computer:

```
$ python -m plumbery -v
```

You can compare this information with reference information posted at [Plumbery package at PiPy](#). If you have used `pip` to install the software, then you can use it again to upgrade the package:

```
$ sudo pip install --upgrade plumbery
```

Remove the plumbery package

Why would you bother about a small set of files at a computer? Anyway, if needed here is the command to remove Plumbery from a python environment:

```
$ sudo pip uninstall plumbery
```

4.2 Installing plumbery from GitHub

Maybe you want to upgrade the on-line documentation of Plumbery. For this you have to edit the *.rst* files that are in the *docs* directory. Or you want to extend plumbery engine, for a new or special usage. Or you would like to troubleshoot an issue and put some *print()* statements in the code. And why not fix a bug or even implement a new feature?

In all these situations, you would like to get a full copy of all files, and change them at will on your own computer.

4.2.1 Install the plumbery development environment

Some precautions are needed if you really want to contribute to the Plumbery project. This is not really difficult, if you follow comprehensive instructions provided at [How to contribute to Plumbery?](#)

On the other hand, to dump a full copy of the software then you can clone the latest development version from [Plumbery repository at GitHub](#):

```
$ sudo apt-get install python-pip python-dev git
$ sudo pip install -e git+https://github.com/bernard357/plumbery.git#egg=plumbery
```

4.2.2 Remove the plumbery development environment

Type the following command to clean the python runtime:

```
$ python setup.py develop --uninstall
```

Then you have to go back to the directory where plumbery was downloaded, and remove files by yourself.

4.3 Configure a Linux machine, or a Mac

By default Plumbery reads credentials from the environment of the computer where it is running.

If you are running Ubuntu you could do:

```
$ nano ~/.bash_profile
```

and type text like the following:

```
# credentials to access cloud resources from Dimension Data
export MCP_USERNAME='*** your account name here ***'
export MCP_PASSWORD='*** your password here ***'

# password to access nodes remotely
export SHARED_SECRET='*** password to access nodes ***'
```

4.4 Configure a Windows machine

By default Plumbery reads credentials and other secrets from the environment of the computer where it is running.

Download <https://raw.githubusercontent.com/bagder/ca-bundle/master/ca-bundle.crt> into %APPDATA%libcloud

```
[Environment]::SetEnvironmentVariable("MCP_USERNAME", "myusername", "Process")
[Environment]::SetEnvironmentVariable("MCP_PASSWORD", "mypassword!", "Process")
[Environment]::SetEnvironmentVariable("SSL_CERT_FILE", "C:\Users\Anthony\AppData\Roaming\libcloud\ca-
```

4.5 Test the installation

Ok, now you can check that plumbery can be triggered from the command line:

```
$ python -m plumbery -v
```

This will display the version of the plumbery engine that is at your disposal.

4.6 Do a quick deployment

Open your preferred text editor to create a new file named `fittings.yaml`. Put the following content in it, save the file, and close the editor:

```
locationId: EU6

blueprints:

  - myBlueprint:
      domain:
        name: myDC
      ethernet:
        name: myVLAN
        subnet: 10.1.10.0
      nodes:
        - myServer:
            appliance: 'Ubuntu'
```

This is a very limited fittings file, yet it is all you need to deploy a new server in the data centre of Frankfurt in Germany (Europe).

At this stage you are ready to deploy the fittings plan. The most straightforward command:

```
$ python -m plumber fittings.yaml deploy
```

This will display a sequence of messages so that you can monitor what is done, and check that all steps are executed correctly.

If you hit an issue that you cannot explain, then make plumber more verbose with the debug flag:

```
$ python -m plumber fittings.yaml deploy -d
```

In the end, keep in mind that resources deployed by plumber are costing money to someone! Hopefully, there is a simple way to stop the bill:

```
$ python -m plumber fittings.yaml dispose
```

Congratulations! Plumber has been installed and tested successfully!

The structure and content of fittings file

5.1 Keywords used by plumbery

This is the list of keywords that are known by plumbery, and that can be used in a fittings file. You will find some comprehensive example at the bottom of this page.

Keyword	Required	Description
apiHost	yes	A private API endpoint. See <code>fittings.facility</code>
appliance	yes	Name of the image to deploy, e.g., 'Ubuntu'. See <code>fittings.appliance</code>
backup	no	Cloud backup configuration. See <code>fittings.backup</code>
beachhead	no	Advanced networking setting. See <code>fittings.beachhead</code>
blueprints	yes	A collection of blueprints
cloud-config	no	The cloud-config details. See <code>fittings.cloud-config</code>
cpu	no	The number of CPU, or the CPU configuration. See <code>fittings.compute</code>
default	no	The class of some fittings. See <code>fittings.defaults</code>
defaults	no	A map of default values. See <code>fittings.defaults</code>
description	no	One line of text with hashtags, e.g., 'This is #ubuntu #master node'
disks	no	Storage type and size. See <code>fittings.disks</code>
domain	yes	See <code>fittings.domain</code>
ethernet	yes	See <code>fittings.ethernet</code>
glue	no	See <code>fittings.glue</code>
information	no	A list of strings explaining what the fittings does
links	no	A map, including 'documentation' and 'credit' as possible values
listeners	no	A collection of listener objects, describing load balancers settings
locationId	yes	Target data centre, e.g., 'EU6'. See <code>fittings.facility</code>
memory	no	Amount of RAM in gigabytes. See <code>fittings.compute</code>
monitoring	no	Monitoring plan, either essentials or advanced . No default
nodes	yes	A collection of node objects, describing the servers to be deployed
parameters	no	Settings that can be provided externally to plumbery
regionId	no	Identify API endpoint, e.g., 'dd-ap'. See <code>fittings.facility</code>
running	no	If set to always , then plumbery cannot delete the fittings

5.2 Multiple documents in one fittings file

YAML allows for multiple documents to be assembled in one fittings plan. The separation of documents is done with three dashes at the beginning of a line. The first document is reserved for plumbery parameters, default settings, etc. Therefore the description of blueprints starts on the second document:

```
---
information:
  - "NFS client and server at two different data centres"
---
blueprints:
  ...
```

5.3 Deploying in multiple geographies

Since Plumbery processes each document independently, it is really easy to configure a deployment that spans multiple data centres, like in the following example:

```
---
information:
  - "Multi-Geography deployment example"
---
regionId: dd-eu
locationId: EU6
blueprints:
  ...
---
regionId: dd-na
locationId: NA9
blueprints:
  ...
```

5.4 Combining private and public clouds in a deployment

Private MCPs are set using the `apiHost` parameter, you must also include the datacenter ID of the cloud as the `locationId`. You can then include another document(s) with the public cloud fittings:

```
---
information:
  - "Multi-Geography deployment example"
---
apiHost: my-private-cloud.com
locationId: MY1
blueprints:
  ...
---
regionId: dd-na
locationId: NA9
blueprints:
  ...
```

5.5 Complete example

```
---
information:
  - "Let's Chat server, self-hosted chat for private teams"
```

```

parameters:

  locationId:
    information:
      - "the target data centre for this deployment"
    type: locations.list
    default: EU6

links:
  documentation: https://github.com/DimensionDataCBUSydney/plumbery-contrib/tree/master/fittings/col
  credit: https://mborgerson.com/setting-up-lets-chat-on-ubuntu

defaults:

  cloud-config:

    ssh_keys:
      rsa_private: |
        {{ key.rsa_private }}
      rsa_public: "{{ key.rsa_public }}"

    users:
      - default

      - name: ubuntu
        sudo: 'ALL=(ALL) NOPASSWD:ALL'
        ssh-authorized-keys:
          - "{{ key.rsa_public }}"
          - "{{ local.rsa_public }}"

    disable_root: true
    ssh_pwauth: false

---

locationId: "{{ parameter.locationId }}"

blueprints:

  - letschat:

    domain:
      name: myDomain
      ipv4: 2

    ethernet:
      name: myNetwork
      subnet: 10.0.0.0

    nodes:
      - letschat01:

        description: "#chat server #ubuntu"

        information:
          - "this is the Let's Chat server for our team"
          - "browse http://{{ node.public }}:5000/ to enter conversations"

```

```
appliance: 'Ubuntu 14'
cpu: 8
memory: 32

disks:
  - 1 50 standard

glue:
  - internet 22 5000

monitoring: essentials

cloud-config:
  hostname: "{{ node.name }}"

packages:
  - nodejs
  - npm
  - nodejs-legacy
  - mongodb
  - mongodb-server
  - git

runcmd:

  - echo "==== Growing LVM with added disk"
  - pvcreate /dev/sdb
  - vgextend rootvol00 /dev/sdb
  - lvextend -l +100%FREE /dev/mapper/rootvol00-rootlv00
  - resize2fs /dev/mapper/rootvol00-rootlv00

  - echo "==== Handling ubuntu identity"
  - cp -n /etc/ssh/ssh_host_rsa_key /home/ubuntu/.ssh/id_rsa
  - cp -n /etc/ssh/ssh_host_rsa_key.pub /home/ubuntu/.ssh/id_rsa.pub
  - chown ubuntu:ubuntu /home/ubuntu/.ssh/*

  - echo "==== Installing Let's Chat"
  - cd /home/ubuntu
  - git clone https://github.com/sdelements/lets-chat.git
  - cd lets-chat
  - npm install
  - cp settings.yml.sample settings.yml
  - sed -i "/host:/s/'[^']*'/'{{ node.name }}'/" settings.yml

  - echo "==== Starting the server"
  - npm start
```

How to contribute to Plumbery?

Contributions are welcome, and they are greatly appreciated! Every little bit helps, and credit will always be given.

6.1 You are not a developer? We are glad that you are involved

You can contribute in many ways:

- submit feedback
- report bugs
- write documentation
- fix bugs
- implement features

6.1.1 Submit Feedback

The best way to send feedback is to file an issue. The place to go is [Plumbery issues at GitHub](#). There you can check if your feedback is new, or if you can align with others.

If you are proposing a feature:

- Explain in detail how it would work.
- Keep the scope as narrow as possible, to make it easier to implement.
- Remember that this is a volunteer-driven project, and that contributions are welcome :)

6.1.2 Report Bugs

Have you identified some bug? Thanks to go to [Plumbery issues at GitHub](#). This is the place where issues are documented, discussed, and fixed. We really value your time and effort to report bugs.

If you are reporting a bug, please include:

- Your operating system name and version.
- Any details about your local setup that might be helpful in troubleshooting.
- Detailed steps to reproduce the bug.

6.1.3 Write Documentation

Plumbery could always use more documentation, whether as part of the official plumbery docs, in docstrings, or even on the web in blog posts, articles, and such.

6.1.4 Fix Bugs

Look through the GitHub issues for bugs. Anything tagged with “bug” is open to whoever wants to implement it.

6.1.5 Implement Features

Look through the GitHub issues for features. Anything tagged with “feature” is open to whoever wants to implement it.

6.2 Ready to contribute? Here’s how to set up Plumbery for local development

1. Fork the *plumbery* repo on [GitHub](#). If you do not have an account there yet, you have to create one, really. This is provided for free, and will make you a proud member of a global community that matters. Once you have authenticated, visit the [Plumbery repository at GitHub](#) and click on the *Fork* link.
2. Clone your fork locally:

```
$ git clone git@github.com:your_name_here/plumbery.git
```

3. Install your local copy into a virtualenv. Assuming you have `virtualenvwrapper` installed, this is how you set up your fork for local development:

```
$ mkvirtualenv plumbery
$ cd plumbery/
$ python setup.py develop
```

4. Create a branch for local development:

```
$ git checkout -b name-of-your-bugfix-or-feature
```

Now you can make your changes locally.

5. When you’re done making changes, check that your changes pass `flake8` and the tests:

```
$ make lint
$ make test
$ make coverage
```

6. Commit your changes and push your branch to GitHub:

```
$ git add .
$ git commit -m "Your detailed description of your changes."
$ git push origin name-of-your-bugfix-or-feature
```

7. Submit a pull request through the GitHub website.

6.3 Pull Request Guidelines

Before you submit a pull request, check that it meets these guidelines:

1. The pull request should include tests.
2. If the pull request adds functionality, the docs should be updated. Put your new functionality into a function with a docstring, and add the feature to the list in `README.rst`.
3. Check [Plumbery continuous integration at Travis CI](#) and make sure that the tests pass there.

History

7.1 1.0.0 (2016-07-15)

- Use libcloud 1.1.0
- Parameters are now prepended with the keyword *parameter* instead of appended (Breaking change)
- Documentation improvements

7.2 0.9.9 (2016-06-10)

- Parse Terraform logging correctly (FIX)

7.3 0.9.8 (2016-06-10)

- Added error logs to multi-clouds

7.4 0.9.7 (2016-06-08)

- Support for Python 3

7.5 0.9.6 (2016-06-06)

- Added multicloud support with Hashicorp terraform, supports parameters being passed to TF vars files directly
- Add colored logging
- Multiple code fixes
- Adapt code for python 3 compatibility
- Passing Travis CI for python 2.7, but not yet for python 3

7.6 0.9.5 (2016-05-28)

- 35 reference configurations are now provided in separate plumbery-contrib
- Passing Travis CI
- 3787 python statements
- Test coverage 53%

7.7 0.9.4 (2016-05-27)

- Adapt Dockerfile to Windows targets
- Fix various bugs for windows - UAT for windows complete.

7.8 0.9.3 (2016-05-27)

- Added pretty logging

7.9 0.9.2 (2016-05-27)

- Fixed bugs around private MCPs

7.10 0.9.1 (2016-05-26)

- Added validation phase to all post-deploy configuration tasks

7.11 0.9.0 (2016-05-25)

- Added initial windows polisher for command execution

7.12 0.8.1 (2016-05-23)

- Remove windows line endings from unix shell fixtures.
- Accept static IPv4 settings
- Support ICMP, TCP, UDP and IP in firewall definitions
- Add alias 'primary' for main NIC network
- Upgrade VMware tools automatically
- Fix calls to API 2.2
- Add the ability to read fittings plan from stdin
- Add entry point to launch plumbery from command line

7.13 0.8.0 (2016-04-13)

- Add the capability to backup nodes

7.14 0.7.0 (2016-04-06)

- Rename 'spit' to 'configure'

7.15 0.6.0 (2016-03-15)

- Added private MCP support (compute only)
- Updated Dockerfile to support parameters
- Set 'finalize' (or 'finalise' for the English) as the new phrase for 'polish'
- Removed coordinate and city lookups (we are adding 40 MCPs this year, I really don't want to maintain that list!)

7.16 0.5.0 (2016-03-13)

- Handover project governance to Dimension Data R&D team
- 26 tutorials are now provided in separate plumbery-contribs project at GitHub
- Introduce deployment parameters (location, domain name, network name, ...)
- Add tutorial on log consolidation (ELK)
- Add tutorial on disk orchestration (standard, high-performance, economy)
- Secure SSH communications in most tutorials (no password, no root access)
- Add virtual storage to many tutorials
- Derive regionId from locationId to streamline fittings plan
- Add action 'refresh' to ease iterations in fittings plan
- Add option -p to load parameters from the command line
- Accept web links for fittings plan and for parameters file
- Restructure code to accomodate for more complex data processing
- Fix multiple bugs and errors
- Augment test coverage
- Improve information provided in debug mode
- Passing Travis CI
- 3286 python statements
- Test coverage 53%

7.17 0.4.3 (2016-02-28)

- A total of 24 various tutorials is provided
- Enhance tutorials on Docker and Kubernetes with virtual storage, SSH keys, class-based definition, and updated tutorial
- Add tutorial on iTop, an open source solution for ITIL management
- Add the management of virtual disks, including tiered storage
- Add tutorial showing partitioning, formatting and mounting of virtual disks
- Enhance the usage documentation page
- Accept pseudo-target of blueprints
- Add classes of nodes to streamline large deployments
- Include the library of Customer Images
- Normalize information reported by the ‘inventory’ command
- Expose geolocalisation information to nodes if needed
- Generalize the usage of ‘information:’ to support active documentation of fittings plans
- Put fittings plan in context with the new ‘links:’ directive
- Reduce requirements in fittings plan passed as a Python dictionary
- Add power off as secondary mechanism to stop nodes, when graceful shutdown is not enough
- Allow for CPU and memory dynamic reconfiguration
- Add option -s to run plumbery in safe mode from the command line
- Report on time spent in the cloud while working
- Fix multiple bugs and errors
- Augment test coverage
- Improve information provided in debug mode
- Passing Travis CI
- 3121 python statements
- Test coverage 50%

7.18 0.4.2 (2016-02-14)

- Add tutorials on Docker, Kubernetes, Docker Swarm – DevOps
- Add tutorial to create a Stackstorm server – DevOps too
- Add tutorial to create a swarm of pingers with MQTT and Kibana – IOT is coming
- A total of 23 various tutorials is provided
- Enhance the documentation accordingly
- Enhance the documentation page on cloud-config
- Add help for Windows users

- Add the generation of uuid when needed (ceph cluster)
- Add the capability to assign public IPv4 when needed (automatic mode)
- Expand information reported by the ‘inventory’ command
- Now support settings for a proxy for all interactions with API endpoints
- Restructure code and improve performance with better cache
- Fix multiple bugs and errors
- Improve information provided in debug mode
- Passing Travis CI
- 2769 python statements
- Test coverage 48%

7.19 0.4.1 (2016-01-24)

- Provision of an advanced tutorial that demonstrates most interesting features
- A total of 18 various tutorials is provided
- Enhance the documentation accordingly
- Leverage cloud-init with global directives shared by all nodes
- Extend dynamic variables to random, on-demand RSA keys
- Enforce password-free communications, and ipv6 communications
- Add actions ‘wipe’ and ‘secrets’
- Restructure code and consolidate functions
- Fix multiple bugs and errors
- Expand information provided in debug mode
- Passing Travis CI
- 2717 python statements
- Test coverage 47%

7.20 0.4.0 (2016-01-17)

- Add cloud-init for node contextualisation
- Inject dynamic variables (e.g., node addresses) to node contextualisation
- Generate and store random secrets for secured node contextualisation
- Enhance usage page in the documentation
- Add a full pack of tutorials in the on-line documentation
- Add actions ‘deploy’ and ‘dispose’ to streamline usage
- Add polishers ‘ping’ and ‘information’
- Restructure code and consolidate functions

- Fix multiple bugs and errors
- Expand information provided in debug mode
- Passing Travis CI
- Test coverage 47%

7.21 0.3.4 (2016-01-06)

- Add tutorials to the documentation
- Fix some errors
- Expand information provided in debug mode
- Passing Travis CI
- Test coverage 39%

7.22 0.3.3 (2015-12-31)

- Create load balancers with pools of nodes
- Streamline messages in safe mode
- Options to run in debug or in quiet mode
- Load fittings plan as a string
- Limit plumbing to some locations if needed
- Make flake8 as happy as possible (PEP8 enforcement)
- Add software documentation for polishers ansible, inventory and rub
- Split software documentation in multiple pages
- Add a first documented use case to the on-line documentation
- Restructure code of the core engine
- Passing Travis CI
- Test coverage 39%

7.23 0.3.2 (2015-12-23)

- Run plumberly from the command-line
- Accept custom polisher from the command line too
- Release the first comprehensive on-line documentation at ReadTheDocs
- Illustrate new features in one demonstration fittings plan (the Gigafox project)
- Consolidate sample scripts and python programs for demonstrations
- Allow multiple network interfaces per node
- Reserve and manage pools of public IPv4 addresses

- Add address translation to nodes equipped with public IPv4 addresses
- Add firewall rules for nodes directly visible from the internet
- Wait for nodes to be deployed before polishing them
- Add new strategies to communicate with nodes over the network
- Fix the usage of puppet for Apache nodes
- Improve reporting messages
- Enhance code documentation
- Fix behaviour on multiple abnormal conditions
- Passing Travis CI
- Test coverage 40%

7.24 0.2.4 (2015-12-13)

- Extend coverage of the Gigafox project
- Add monitoring to nodes created
- Run polisher ‘spit’ right after build to finalise setup of nodes
- Produce inventory with polisher of same name
- Introduce puppet manifests to polisher ‘rub’
- Add file copy to remote nodes
- Introduce basement blueprints
- Improve reporting messages
- Enhance code documentation
- Fix behaviour on abnormal situations
- Passing Travis CI
- Test coverage 46%

7.25 0.2.3 (2015-12-07)

- Introduce Gigafox project, to demonstrate deployment across multiple places
- Create firewall rules to allow traffic between networks
- Cache information to save on API calls and to accelerate the processing
- Improve the state engine
- Passing Travis CI
- Test coverage 48%

7.26 0.2.2 (2015-12-05)

- Use ipv6 when possible to connect to remote nodes
- Manual tests to improve usage on specific conditions, e.g., against MCP 1.0
- Label expansion to facilitate node handling, e.g., mongo[1..20]
- Allow for destruction of networks and of domain networks
- Passing Travis CI
- Test coverage 55%

7.27 0.2.1 (2015-12-02)

- Code refactoring into a new module for nodes
- fake8 complains only about long lines and line termination
- Passing Travis CI
- Test coverage 59%

7.28 0.2.0 (2015-11-29)

- Allow for node bootstrapping via SSH
- Push SSH public key
- Update Linux packages
- Install Docker
- Provide an inventory of running nodes
- Build inventory for ansible
- Extension mechanism called polishers
- Add demonstration scripts related to polishers
- A lot of docstring has been added
- fake8 complains only about long lines and line termination
- Passing Travis CI
- Test coverage 59%

7.29 0.1.2 (2015-11-27)

- First pull request fully processed
- Docstring for all modules
- Passing Travis CI
- Test coverage 62%

7.30 0.1.0 (2015-11-20)

- First release on PyPI.

plumbery package

8.1 Subpackages

8.1.1 plumbery.actions package

Submodules

plumbery.actions.echo module

Module contents

8.1.2 plumbery.polishers package

Submodules

plumbery.polishers.ansible module

plumbery.polishers.configure module

plumbery.polishers.information module

plumbery.polishers.inventory module

plumbery.polishers.ping module

plumbery.polishers.prepare module

Module contents

8.2 Submodules

8.2.1 plumbery.action module

8.2.2 plumbery.bootstrap module

plumbery.bootstrap.**download_file** (*url*, *output_dir*)

`plumberly.bootstrap.download_manifest` (*url*, *output_dir*)

`plumberly.bootstrap.main` (*args*)

`plumberly.bootstrap.parse_args` (*args*)

8.2.3 plumberly.engine module

8.2.4 plumberly.exception module

exception `plumberly.exception.ConfigurationError` (*message*)

Bases: `plumberly.exception.PlumberlyException`

exception `plumberly.exception.PlumberlyException` (*message*)

Bases: `exceptions.Exception`

8.2.5 plumberly.facility module

8.2.6 plumberly.infrastructure module

8.2.7 plumberly.nodes module

8.2.8 plumberly.plogging module

8.2.9 plumberly.polisher module

8.2.10 plumberly.terraform module

8.2.11 plumberly.text module

8.2.12 plumberly.util module

8.3 Module contents

Are you looking for a cloud plumber? We hope this one will be useful to you

- [genindex](#)

p

plumbery, 34
plumbery.actions, 33
plumbery.bootstrap, 33
plumbery.exception, 34
plumbery.polishers, 33

C

ConfigurationError, 34

D

download_file() (in module plumbery.bootstrap), 33

download_manifest() (in module plumbery.bootstrap), 33

M

main() (in module plumbery.bootstrap), 34

P

parse_args() (in module plumbery.bootstrap), 34

plumbery (module), 34

plumbery.actions (module), 33

plumbery.bootstrap (module), 33

plumbery.exception (module), 34

plumbery.polishers (module), 33

PlumberyException, 34