
Cloud-Init

Release 0.7.8

September 30, 2016

1	Summary	3
1.1	Capabilities	3
1.2	Availability	3
1.3	Formats	4
1.4	Directory layout	7
1.5	Cloud config examples	8
1.6	Datasources	20
1.7	Modules	29
1.8	Merging User-Data Sections	36
1.9	More information	39
1.10	Hacking on cloud-init	39
	Python Module Index	41

Everything about cloud-init, a set of python scripts and utilities to make your cloud images be all they can be!

Summary

Cloud-init is the *defacto* multi-distribution package that handles early initialization of a cloud instance.

1.1 Capabilities

- Setting a default locale
- Setting a instance hostname
- Generating instance ssh private keys
- Adding ssh keys to a users `.ssh/authorized_keys` so they can log in
- Setting up ephemeral mount points

1.1.1 User configurability

Cloud-init 's behavior can be configured via user-data.

User-data can be given by the user at instance launch time.

This is done via the `--user-data` or `--user-data-file` argument to `ec2-run-instances` for example.

- Check your local clients documentation for how to provide a *user-data* string or *user-data* file for usage by cloud-init on instance creation.

1.2 Availability

It is currently installed in the [Ubuntu Cloud Images](#) and also in the official [Ubuntu](#) images available on EC2.

Versions for other systems can be (or have been) created for the following distributions:

- Ubuntu
- Fedora
- Debian
- RHEL
- CentOS

- *and more...*

So ask your distribution provider where you can obtain an image with it built-in if one is not already available

1.3 Formats

User data that will be acted upon by cloud-init must be in one of the following types.

1.3.1 Gzip Compressed Content

Content found to be gzip compressed will be uncompressed. The uncompressed data will then be used as if it were not compressed. This is typically useful because user-data is limited to ~16384¹ bytes.

1.3.2 Mime Multi Part Archive

This list of rules is applied to each part of this multi-part file. Using a mime-multi part file, the user can specify more than one type of data.

For example, both a user data script and a cloud-config type could be specified.

Supported content-types:

- text/x-include-once-url
- text/x-include-url
- text/cloud-config-archive
- text/upstart-job
- text/cloud-config
- text/part-handler
- text/x-shellscript
- text/cloud-boothook

Helper script to generate mime messages

```
#!/usr/bin/python

import sys

from email.mime.multipart import MIMEMultipart
from email.mime.text import MIMEText

if len(sys.argv) == 1:
    print("%s input-file:type ..." % (sys.argv[0]))
    sys.exit(1)

combined_message = MIMEMultipart()
for i in sys.argv[1:]:
    (filename, format_type) = i.split(":", 1)
    with open(filename) as fh:
```

¹ See your cloud provider for applicable user-data size limitations...

```

        contents = fh.read()
        sub_message = MIMEText(contents, format_type, sys.getdefaultencoding())
        sub_message.add_header('Content-Disposition', 'attachment; filename="%s"' % (filename))
        combined_message.attach(sub_message)

print(combined_message)

```

1.3.3 User-Data Script

Typically used by those who just want to execute a shell script.

Begins with: `#!` or `Content-Type: text/x-shellscript` when using a MIME archive.

Example

```

$ cat myscript.sh

#!/bin/sh
echo "Hello World. The time is now $(date -R)!" | tee /root/output.txt

$ euca-run-instances --key mykey --user-data-file myscript.sh ami-a07d95c9

```

1.3.4 Include File

This content is a `include` file.

The file contains a list of urls, one per line. Each of the URLs will be read, and their content will be passed through this same set of rules. Ie, the content read from the URL can be gzipped, mime-multi-part, or plain text.

Begins with: `#include` or `Content-Type: text/x-include-url` when using a MIME archive.

1.3.5 Cloud Config Data

Cloud-config is the simplest way to accomplish some things via user-data. Using cloud-config syntax, the user can specify certain things in a human friendly format.

These things include:

- apt upgrade should be run on first boot
- a different apt mirror should be used
- additional apt sources should be added
- certain ssh keys should be imported
- *and many more...*

Note: The file must be valid yaml syntax.

See the [Cloud config examples](#) section for a commented set of examples of supported cloud config formats.

Begins with: `#cloud-config` or `Content-Type: text/cloud-config` when using a MIME archive.

1.3.6 Upstart Job

Content is placed into a file in `/etc/init`, and will be consumed by upstart as any other upstart job.

Begins with: `#upstart-job` or `Content-Type: text/upstart-job` when using a MIME archive.

1.3.7 Cloud Boothook

This content is boothook data. It is stored in a file under `/var/lib/cloud` and then executed immediately. This is the earliest hook available. Note, that there is no mechanism provided for running only once. The boothook must take care of this itself. It is provided with the instance id in the environment variable `INSTANCE_I`. This could be made use of to provide a 'once-per-instance' type of functionality.

Begins with: `#cloud-boothook` or `Content-Type: text/cloud-boothook` when using a MIME archive.

1.3.8 Part Handler

This is a part-handler. It will be written to a file in `/var/lib/cloud/data` based on its filename (which is generated). This must be python code that contains a `list_types` method and a `handle_type` method. Once the section is read the `list_types` method will be called. It must return a list of mime-types that this part-handler handles.

The `handle_type` method must be like:

```
def handle_part(data, ctype, filename, payload):
    # data = the cloudfinit object
    # ctype = "__begin__", "__end__", or the mime-type of the part that is being handled.
    # filename = the filename of the part (or a generated filename if none is present in mime data)
    # payload = the parts' content
```

Cloud-init will then call the `handle_type` method once at begin, once per part received, and once at end. The begin and end calls are to allow the part handler to do initialization or teardown.

Begins with: `#part-handler` or `Content-Type: text/part-handler` when using a MIME archive.

Example

```
1 #part-handler
2 # vi: syntax=python ts=4
3
4 def list_types():
5     # return a list of mime-types that are handled by this module
6     return(["text/plain", "text/go-cubs-go"])
7
8 def handle_part(data, ctype, filename, payload):
9     # data: the cloudfinit object
10    # ctype: '__begin__', '__end__', or the specific mime-type of the part
11    # filename: the filename for the part, or dynamically generated part if
12    #           no filename is given attribute is present
13    # payload: the content of the part (empty for begin or end)
14    if ctype == "__begin__":
15        print "my handler is beginning"
16        return
17    if ctype == "__end__":
18        print "my handler is ending"
```

```

19     return
20
21     print "==== received ctype=%s filename=%s ==== " % (ctype, filename)
22     print payload
23     print "==== end ctype=%s filename=%s" % (ctype, filename)

```

Also this [blog post](#) offers another example for more advanced usage.

1.4 Directory layout

Cloudinit's directory structure is somewhat different from a regular application:

```

/var/lib/cloud/
- data/
  - instance-id
  - previous-instance-id
  - datasource
  - previous-datasource
  - previous-hostname
- handlers/
- instance
- instances/
  i-00000XYZ/
    - boot-finished
    - cloud-config.txt
    - datasource
    - handlers/
    - obj.pkl
    - scripts/
    - sem/
    - user-data.txt
    - user-data.txt.i
- scripts/
  - per-boot/
  - per-instance/
  - per-once/
- seed/
- sem/

```

`/var/lib/cloud`

The main directory containing the cloud-init specific subdirectories. It is typically located at `/var/lib` but there are certain configuration scenarios where this can be altered.

TBD, describe this overriding more.

`data/`

Contains information related to instance ids, datasources and hostnames of the previous and current instance if they are different. These can be examined as needed to determine any information related to a previous boot (if applicable).

`handlers/`

Custom `part-handlers` code is written out here. Files that end up here are written out with in the scheme of `part-handler-XYZ` where XYZ is the handler number (the first handler found starts at 0).

`instance`

A symlink to the current `instances/` subdirectory that points to the currently active instance (which is active is dependent on the datasource loaded).

`instances/`

All instances that were created using this image end up with instance identifier subdirectories (and corresponding data for each instance). The currently active instance will be symlinked the the `instance` symlink file defined previously.

`scripts/`

Scripts that are downloaded/created by the corresponding `part-handler` will end up in one of these subdirectories.

`seed/`

TBD

`sem/`

Cloud-init has a concept of a module semaphore, which basically consists of the module name and its frequency. These files are used to ensure a module is only ran *per-once*, *per-instance*, *per-always*. This folder contains semaphore *files* which are only supposed to run *per-once* (not tied to the instance id).

1.5 Cloud config examples

1.5.1 Including users and groups

```
1 # Add groups to the system
2 # The following example adds the ubuntu group with members foo and bar and
3 # the group cloud-users.
4 groups:
5   - ubuntu: [foo,bar]
6   - cloud-users
7
8 # Add users to the system. Users are added after groups are added.
9 users:
10  - default
11  - name: foobar
12    gecos: Foo B. Bar
13    primary-group: foobar
14    groups: users
15    selinux-user: staff_u
16    expiredate: 2012-09-01
17    ssh-import-id: foobar
18    lock_passwd: false
19    passwd: $6$j212wezy$7H/1LT4f9/N3wpgNunhsIqtMj62OKiS3nyNwuzouQc3u7MbYCarYeAHWYPYb2FT.lbioDm2RrkJ
20  - name: barfoo
21    gecos: Bar B. Foo
22    sudo: ALL=(ALL) NOPASSWD:ALL
23    groups: users, admin
24    ssh-import-id: None
25    lock_passwd: true
26    ssh-authorized-keys:
27      - <ssh pub key 1>
28      - <ssh pub key 2>
29  - name: cloudy
30    gecos: Magic Cloud App Daemon User
```

```

31 inactive: true
32 system: true
33
34 # Valid Values:
35 #   name: The user's login name
36 #   gecos: The user name's real name, i.e. "Bob B. Smith"
37 #   homedir: Optional. Set to the local path you want to use. Defaults to
38 #             /home/<username>
39 #   primary-group: define the primary group. Defaults to a new group created
40 #                 named after the user.
41 #   groups: Optional. Additional groups to add the user to. Defaults to none
42 #   selinux-user: Optional. The SELinux user for the user's login, such as
43 #                 "staff_u". When this is omitted the system will select the default
44 #                 SELinux user.
45 #   lock_passwd: Defaults to true. Lock the password to disable password login
46 #   inactive: Create the user as inactive
47 #   passwd: The hash -- not the password itself -- of the password you want
48 #            to use for this user. You can generate a safe hash via:
49 #            mkpasswd --method=SHA-512 --rounds=4096
50 #            (the above command would create from stdin an SHA-512 password hash
51 #            with 4096 salt rounds)
52 #
53 #           Please note: while the use of a hashed password is better than
54 #           plain text, the use of this feature is not ideal. Also,
55 #           using a high number of salting rounds will help, but it should
56 #           not be relied upon.
57 #
58 #           To highlight this risk, running John the Ripper against the
59 #           example hash above, with a readily available wordlist, revealed
60 #           the true password in 12 seconds on a i7-2620QM.
61 #
62 #           In other words, this feature is a potential security risk and is
63 #           provided for your convenience only. If you do not fully trust the
64 #           medium over which your cloud-config will be transmitted, then you
65 #           should use SSH authentication only.
66 #
67 #           You have thus been warned.
68 #   no-create-home: When set to true, do not create home directory.
69 #   no-user-group: When set to true, do not create a group named after the user.
70 #   no-log-init: When set to true, do not initialize lastlog and faillog database.
71 #   ssh-import-id: Optional. Import SSH ids
72 #   ssh-authorized-keys: Optional. [list] Add keys to user's authorized keys file
73 #   sudo: Defaults to none. Set to the sudo string you want to use, i.e.
74 #         ALL=(ALL) NOPASSWD:ALL. To add multiple rules, use the following
75 #         format.
76 #         sudo:
77 #             - ALL=(ALL) NOPASSWD:/bin/mysql
78 #             - ALL=(ALL) ALL
79 #         Note: Please double check your syntax and make sure it is valid.
80 #         cloud-init does not parse/check the syntax of the sudo
81 #         directive.
82 #   system: Create the user as a system user. This means no home directory.
83 #
84
85 # Default user creation:
86 #
87 # Unless you define users, you will get a 'ubuntu' user on ubuntu systems with the
88 # legacy permission (no password sudo, locked user, etc). If however, you want

```

```

89 # to have the 'ubuntu' user in addition to other users, you need to instruct
90 # cloud-init that you also want the default user. To do this use the following
91 # syntax:
92 #   users:
93 #     - default
94 #     - bob
95 #     - ....
96 # foobar: ...
97 #
98 # users[0] (the first user in users) overrides the user directive.
99 #
100 # The 'default' user above references the distro's config:
101 # system_info:
102 #   default_user:
103 #     name: Ubuntu
104 #     plain_text_passwd: 'ubuntu'
105 #     home: /home/ubuntu
106 #     shell: /bin/bash
107 #     lock_passwd: True
108 #     gecos: Ubuntu
109 #     groups: [adm, audio, cdrom, dialout, floppy, video, plugdev, dip, netdev]

```

1.5.2 Writing out arbitrary files

```

1 #cloud-config
2 # vim: syntax=yaml
3 #
4 # This is the configuration syntax that the write_files module
5 # will know how to understand. encoding can be given b64 or gzip or (gz+b64).
6 # The content will be decoded accordingly and then written to the path that is
7 # provided.
8 #
9 # Note: Content strings here are truncated for example purposes.
10 write_files:
11 -   encoding: b64
12     content: CiMgVGhpcyBmaWxlIGNvbnRyb2xzIHRobzZSBzdGF0ZSBvZiBTRUxpbmV4...
13     owner: root:root
14     path: /etc/sysconfig/selinux
15     permissions: '0644'
16 -   content: |
17         # My new /etc/sysconfig/samba file
18
19         SMBDOPTIONS="-D"
20     path: /etc/sysconfig/samba
21 -   content: !!binary |
22         f0VMRgIBAQAAAAAAAAAAAAAAAAIAPgABAAAAwARAAAAAAAAABAAAAAAAAAAJAVAAAAAAAAAAAAAAAAEAAOAAI
23         AEAHqAdAAYAAAAFAAAAQAAAAAAAAABAAEAAAAAAAAEAAQAAAAAAAAwAEAAAAAAAAADAAQAAAAAAAAAgA
24         AAAAAAAAAAwAAAAQAAAAAAgAAAAAAAAACQAAAAAAAAAJAAAAAAAAcAAAAAAAAABwAAAAAAAAAAQAA
25         ....
26     path: /bin/arch
27     permissions: '0555'
28 -   encoding: gzip
29     content: !!binary |
30         H4sIAIDb/U8C/1NW1E/KzNMvzuBKTC7IV8hIzcnJVYjPL8pJ4QIA6N+MVxsAAAA=
31     path: /usr/bin/hello
32     permissions: '0755'

```

1.5.3 Adding a yum repository

```

1 #cloud-config
2 # vim: syntax=yaml
3 #
4 # Add yum repository configuration to the system
5 #
6 # The following example adds the file /etc/yum.repos.d/epel_testing.repo
7 # which can then subsequently be used by yum for later operations.
8 yum_repos:
9   # The name of the repository
10  epel-testing:
11    # Any repository configuration options
12    # See: man yum.conf
13    #
14    # This one is required!
15    baseurl: http://download.fedoraproject.org/pub/epel/testing/5/$basearch
16    enabled: false
17    failovermethod: priority
18    gpgcheck: true
19    gpgkey: file:///etc/pki/rpm-gpg/RPM-GPG-KEY-EPEL
20    name: Extra Packages for Enterprise Linux 5 - Testing

```

1.5.4 Configure an instances trusted CA certificates

```

1 #cloud-config
2 #
3 # This is an example file to configure an instance's trusted CA certificates
4 # system-wide for SSL/TLS trust establishment when the instance boots for the
5 # first time.
6 #
7 # Make sure that this file is valid yaml before starting instances.
8 # It should be passed as user-data when starting the instance.
9
10 ca-certs:
11   # If present and set to True, the 'remove-defaults' parameter will remove
12   # all the default trusted CA certificates that are normally shipped with
13   # Ubuntu.
14   # This is mainly for paranoid admins - most users will not need this
15   # functionality.
16   remove-defaults: true
17
18   # If present, the 'trusted' parameter should contain a certificate (or list
19   # of certificates) to add to the system as trusted CA certificates.
20   # Pay close attention to the YAML multiline list syntax. The example shown
21   # here is for a list of multiline certificates.
22   trusted:
23     - |
24       -----BEGIN CERTIFICATE-----
25       YOUR-ORGS-TRUSTED-CA-CERT-HERE
26       -----END CERTIFICATE-----
27     - |
28       -----BEGIN CERTIFICATE-----
29       YOUR-ORGS-TRUSTED-CA-CERT-HERE
30       -----END CERTIFICATE-----

```

1.5.5 Configure an instances resolv.conf

Note: when using a config drive and a RHEL like system resolv.conf will also be managed ‘automatically’ due to the available information provided for dns servers in the config drive network format. For those that wish to have different settings use this module.

```

1 #cloud-config
2 #
3 # This is an example file to automatically configure resolv.conf when the
4 # instance boots for the first time.
5 #
6 # Ensure that your yaml is valid and pass this as user-data when starting
7 # the instance. Also be sure that your cloud.cfg file includes this
8 # configuration module in the appropriate section.
9 #
10 manage-resolv-conf: true
11
12 resolv_conf:
13   nameservers: ['8.8.4.4', '8.8.8.8']
14   searchdomains:
15     - foo.example.com
16     - bar.example.com
17   domain: example.com
18   options:
19     rotate: true
20     timeout: 1

```

1.5.6 Install and run chef recipes

```

1 #cloud-config
2 #
3 # This is an example file to automatically install chef-client and run a
4 # list of recipes when the instance boots for the first time.
5 # Make sure that this file is valid yaml before starting instances.
6 # It should be passed as user-data when starting the instance.
7 #
8 # This example assumes the instance is 12.04 (precise)
9 #
10 # The default is to install from packages.
11
12 # Key from http://apt.opscode.com/packages@opscode.com.gpg.key
13 apt:
14   sources:
15     - source: "deb http://apt.opscode.com/ $RELEASE-0.10 main"
16     key: |
17       -----BEGIN PGP PUBLIC KEY BLOCK-----
18       Version: GnuPG v1.4.9 (GNU/Linux)
19
20       mQGIBepPc7QRBADfsOkZU6KZK+YmKw4wev5mjKJEkVGlus+NxW8wItX5sGa6kdUu
21       twAyj7Yr92rF+ICFEP3gGU6+lGo0Nve7KxkN/1W7/m3G4zuk+ccIKmjp8KS3qn99
22       dxy64vcji9jI11Va+XXOGIp0G8GEaj7mbkixL/bMeGfdMlv8Gf2XPPP9vwCgn/GC
23       JKacfnw7MpLKUHOYS1b//JsEAJqao3ViNfav83jJKEkD8cf59Y8xKia5OpZqTK5W
24       ShVnNWS3U5IVQk10ZDH97Qn/YrK387H4CyhLE9mxPXs/ull8ioiaars/q2MEKU2I
25       XKfV21eMLO9LYd6Ny/Kqj8o5WQK2J6+NAhSwvthZcIEphcFignIuobP+B5wNFQpe
26       DbKfA/0WvN2OwFeWRcmmd3Hz7nHTpcnSF+4QX6yHRF/5BgxkG6IqBIACQbzPn6Hm
27       sMtm/SVf1lizmDqSsQptCrOZILfLX/mE+YO1+CwWSHh1+YsFts1WOUh1EhQD26aO
28

```

```

29     Z84HuHV5HFRWjDLw9LriltBVQcXbpfSrRP5bdr7Wh8vqhJTPjrQnT3BzY29kZSBQ
30     YWNrYWdlcyA8cGFja2FnZXNAb3BzY29kZS5jb20+iGAEEeECACAFakppC7QCGwMG
31     CwkIBwMCBBUCCAMEFgIDAQIeAQIXgAAKCRApQKupg++Ca j8sAKCOXmdG36gWji/K
32     +o+XtBfvdMnFYQCfTCEWxRy2BnzLoBBFCjDSK6sJqCu5Ag0ESmkLtBAIAIO2SwlR
33     lU5i6gTop42RHWW7/pmW78CwUqJnYqnXROrt3h9F9xrsGkH0Fh1FRtsnncgzIhvh
34     DLQnRHnkXm0ws0jV0PF74ttoUT6BLAUsFi2SPP1zYNJ9H9fhhK/pji jtAcQwdgxu
35     wwNJ5xCEscBZCjhSRXm0d30bKl049Cow8ZiBhtnXVP41c9QWOzX/LaGZsKQZnaMx
36     EzDk8dyyctR2f03vRSVyTFGgdpUcpbr9eTFVgikCa6ODEBv+0BnCH6yGTxwBid9g
37     w0ole/2DviKUWCC+AlAUOubLmOIGFBuI4UR+rux9affbHcLIOTiKQXv79lW3P7W8
38     AAfniSQKfPWXrrcAAwUH/2XBqD4Uxhbs25HDUUIm/m6Gnlj6EsStg8n0nMggLhuN
39     QmPfoNByMPUqvA7sULyfr6xCYzbzRNxABHSpf85FzGQ29RF4xsA4vOOU8RDIYQ9X
40     Q8NqqR6pydprRFqWe47hsAN7BoYuhWqTtOLSBmnAnzTR5pUROqcquWYiiEavZixJ
41     3ZRAq/HMGioJEtMFrsvsZjGXuze7f0ytfrlZYeLVWnL9Bd32CueBli7dhYwkFe+V
42     Ep5jWOCj02ClwHcwt+uIRDJV6TdtbIiBYAdOMPk15+VBdweBXwMuYXr76+A7VeDL
43     zIhi7tKf06WiwjKZq0dzctsJJjtIfr4K4vbiD90jg1iISQQYEQIACQUCSmkLtAIb
44     DAAKCRApQKupg++CauISAJ9CxyPOKhOxalBnVTLeNUKAHGg2gACeIsbobaD4ZHG
45     0GLl8EkfA8uhluM=
46     =zKAm
47     -----END PGP PUBLIC KEY BLOCK-----
48

```

```

49 chef:

```

```

50
51 # Valid values are 'gems' and 'packages' and 'omnibus'
52 install_type: "packages"
53
54 # Boolean: run 'install_type' code even if chef-client
55 #     appears already installed.
56 force_install: false
57
58 # Chef settings
59 server_url: "https://chef.yourorg.com:4000"
60
61 # Node Name
62 # Defaults to the instance-id if not present
63 node_name: "your-node-name"
64
65 # Environment
66 # Defaults to '_default' if not present
67 environment: "production"
68
69 # Default validation name is chef-validator
70 validation_name: "yourorg-validator"
71 # if validation_cert's value is "system" then it is expected
72 # that the file already exists on the system.
73 validation_cert: |
74     -----BEGIN RSA PRIVATE KEY-----
75     YOUR-ORGS-VALIDATION-KEY-HERE
76     -----END RSA PRIVATE KEY-----
77
78 # A run list for a first boot json
79 run_list:
80   - "recipe[apache2]"
81   - "role[db]"
82
83 # Specify a list of initial attributes used by the cookbooks
84 initial_attributes:
85   apache:
86     prefork:

```



```

44 1MASW2CfaEiOEi9F1R3R4Q1z4ix+iNoHiUDTjazw/tZwEdxaQXQVLwgTGRwVa+aA
45 qbutJKi93MILLwIDAQABo3kwdzA4BglghkgBhvhCAQ0EKxYpUHVwcGV0IFJlYnkv
46 T3B1blNTTCBHZW51cmF0ZWQgQ2VydGlmawNhdGUwDwYDVR0TAQH/BAUwAwEB/zAd
47 BgNVHQ4EFgQUu4+jHB+GYE5Vxo+ol1OAhevspjAwCwYDVR0PBAQDAgEGMA0GCSqG
48 S1b3DQEBBQUAA4GBAH/rx1UIjwNb3n7TXJcDJ6MMHULwjr03BDJXKb34U1ndkpaf
49 +GalzPXWa7bO908M9I8RnPfvtKnteLbvgTK+h+zX1XCty+S2EQWk29i2AdoqOTxb
50 hppiGMp0tT5Havu4aceCXiy2crVcudj3NFciy8X66SoECemW9UYDCb9T5D0d
51 -----END CERTIFICATE-----

```

1.5.8 Add apt repositories

```

1 #cloud-config
2
3 # Add apt repositories
4 #
5 # Default: auto select based on cloud metadata
6 # in ec2, the default is <region>.archive.ubuntu.com
7 # apt:
8 #   primary:
9 #     - arches [default]
10 #     uri:
11 #     use the provided mirror
12 #     search:
13 #     search the list for the first mirror.
14 #     this is currently very limited, only verifying that
15 #     the mirror is dns resolvble or an IP address
16 #
17 # if neither mirror is set (the default)
18 # then use the mirror provided by the DataSource found.
19 # In EC2, that means using <region>.ec2.archive.ubuntu.com
20 #
21 # if no mirror is provided by the DataSource, but 'search_dns' is
22 # true, then search for dns names '<distro>-mirror' in each of
23 # - fqdn of this host per cloud metadata
24 # - localdomain
25 # - no domain (which would search domains listed in /etc/resolv.conf)
26 # If there is a dns entry for <distro>-mirror, then it is assumed that there
27 # is a distro mirror at http://<distro>-mirror.<domain>/<distro>
28 #
29 # That gives the cloud provider the opportunity to set mirrors of a distro
30 # up and expose them only by creating dns entries.
31 #
32 # if none of that is found, then the default distro mirror is used
33 apt:
34   primary:
35     - arches: [default]
36     uri: http://us.archive.ubuntu.com/ubuntu/
37 # or
38 apt:
39   primary:
40     - arches: [default]
41     search:
42     - http://local-mirror.mydomain
43     - http://archive.ubuntu.com
44 # or
45 apt:
46   primary:

```

```
47 - arches: [default]
48   search_dns: True
```

1.5.9 Run commands on first boot

```
1 #cloud-config
2
3 # boot commands
4 # default: none
5 # this is very similar to runcmd, but commands run very early
6 # in the boot process, only slightly after a 'boothook' would run.
7 # bootcmd should really only be used for things that could not be
8 # done later in the boot process. bootcmd is very much like
9 # boothook, but possibly with more friendly.
10 # - bootcmd will run on every boot
11 # - the INSTANCE_ID variable will be set to the current instance id.
12 # - you can use 'cloud-init-boot-per' command to help only run once
13 bootcmd:
14 - echo 192.168.1.130 us.archive.ubuntu.com > /etc/hosts
15 - [ cloud-init-per, once, mymkfs, mkfs, /dev/vdb ]
```

```
1 #cloud-config
2
3 # run commands
4 # default: none
5 # runcmd contains a list of either lists or a string
6 # each item will be executed in order at rc.local like level with
7 # output to the console
8 # - runcmd only runs during the first boot
9 # - if the item is a list, the items will be properly executed as if
10 #   passed to execve(3) (with the first arg as the command).
11 # - if the item is a string, it will be simply written to the file and
12 #   will be interpreted by 'sh'
13 #
14 # Note, that the list has to be proper yaml, so you have to quote
15 # any characters yaml would eat (':' can be problematic)
16 runcmd:
17 - [ ls, -l, / ]
18 - [ sh, -xc, "echo $(date) ': hello world!'" ]
19 - [ sh, -c, echo "====hello world'====" ]
20 - ls -l /root
21 - [ wget, "http://slashdot.org", -O, /tmp/index.html ]
```

1.5.10 Alter the completion message

```
1 #cloud-config
2
3 # final_message
4 # default: cloud-init boot finished at $TIMESTAMP. Up $UPTIME seconds
5 # this message is written by cloud-final when the system is finished
6 # its first boot
7 final_message: "The system is finally up, after $UPTIME seconds"
```

1.5.11 Install arbitrary packages

```

1 #cloud-config
2
3 # Install additional packages on first boot
4 #
5 # Default: none
6 #
7 # if packages are specified, this apt_update will be set to true
8 #
9 # packages may be supplied as a single package name or as a list
10 # with the format [<package>, <version>] wherein the specific
11 # package version will be installed.
12 packages:
13 - pwgen
14 - pastebinit
15 - [libpython2.7, 2.7.3-0ubuntu3.1]

```

1.5.12 Run apt or yum upgrade

```

1 #cloud-config
2
3 # Upgrade the instance on first boot
4 # (ie run apt-get upgrade)
5 #
6 # Default: false
7 # Aliases: apt_upgrade
8 package_upgrade: true

```

1.5.13 Adjust mount points mounted

```

1 #cloud-config
2
3 # set up mount points
4 # 'mounts' contains a list of lists
5 # the inner list are entries for an /etc/fstab line
6 # ie : [ fs_spec, fs_file, fs_vfstype, fs_mntops, fs_freq, fs_passno ]
7 #
8 # default:
9 # mounts:
10 # - [ ephemeral0, /mnt ]
11 # - [ swap, none, swap, sw, 0, 0 ]
12 #
13 # in order to remove a previously listed mount (ie, one from defaults)
14 # list only the fs_spec. For example, to override the default, of
15 # mounting swap:
16 # - [ swap ]
17 # or
18 # - [ swap, null ]
19 #
20 # - if a device does not exist at the time, an entry will still be
21 # written to /etc/fstab.
22 # - '/dev' can be omitted for device names that begin with: xvd, sd, hd, vd
23 # - if an entry does not have all 6 fields, they will be filled in
24 # with values from 'mount_default_fields' below.

```

```
25 #
26 # Note, that you should set 'nobootwait' (see man fstab) for volumes that may
27 # not be attached at instance boot (or reboot)
28 #
29 mounts:
30 - [ ephemeral0, /mnt, auto, "defaults,noexec" ]
31 - [ sdc, /opt/data ]
32 - [ xvdh, /opt/data, "auto", "defaults,nobootwait", "0", "0" ]
33 - [ dd, /dev/zero ]
34
35 # mount_default_fields
36 # These values are used to fill in any entries in 'mounts' that are not
37 # complete. This must be an array, and must have 7 fields.
38 mount_default_fields: [ None, None, "auto", "defaults,nobootwait", "0", "2" ]
39
40
41 # swap can also be set up by the 'mounts' module
42 # default is to not create any swap files, because 'size' is set to 0
43 swap:
44   filename: /swap.img
45   size: "auto" # or size in bytes
46   maxsize: size in bytes
```

1.5.14 Call a url when finished

```
1 #cloud-config
2
3 # phone_home: if this dictionary is present, then the phone_home
4 # cloud-config module will post specified data back to the given
5 # url
6 # default: none
7 # phone_home:
8 #   url: http://my.foo.bar/$INSTANCE/
9 #   post: all
10 #   tries: 10
11 #
12 phone_home:
13   url: http://my.example.com/$INSTANCE_ID/
14   post: [ pub_key_dsa, pub_key_rsa, pub_key_ecdsa, instance_id ]
```

1.5.15 Reboot/poweroff when finished

```
1 #cloud-config
2
3 ## poweroff or reboot system after finished
4 # default: none
5 #
6 # power_state can be used to make the system shutdown, reboot or
7 # halt after boot is finished. This same thing can be achieved by
8 # user-data scripts or by runcmd by simply invoking 'shutdown'.
9 #
10 # Doing it this way ensures that cloud-init is entirely finished with
11 # modules that would be executed, and avoids any error/log messages
12 # that may go to the console as a result of system services like
13 # syslog being taken down while cloud-init is running.
```

```

14 #
15 # If you delay '+5' (5 minutes) and have a timeout of
16 # 120 (2 minutes), then the max time until shutdown will be 7 minutes.
17 # cloud-init will invoke 'shutdown +5' after the process finishes, or
18 # when 'timeout' seconds have elapsed.
19 #
20 # delay: form accepted by shutdown. default is 'now'. other format
21 #     accepted is +m (m in minutes)
22 # mode: required. must be one of 'poweroff', 'halt', 'reboot'
23 # message: provided as the message argument to 'shutdown'. default is none.
24 # timeout: the amount of time to give the cloud-init process to finish
25 #         before executing shutdown.
26 # condition: apply state change only if condition is met.
27 #             May be boolean True (always met), or False (never met),
28 #             or a command string or list to be executed.
29 #             command's exit code indicates:
30 #                 0: condition met
31 #                 1: condition not met
32 #             other exit codes will result in 'not met', but are reserved
33 #             for future use.
34 #
35 power_state:
36   delay: "+30"
37   mode: poweroff
38   message: Bye Bye
39   timeout: 30
40   condition: True

```

1.5.16 Configure instances ssh-keys

```

1 #cloud-config
2
3 # add each entry to ~/.ssh/authorized_keys for the configured user or the
4 # first user defined in the user definition directive.
5 ssh_authorized_keys:
6   - ssh-rsa AAAAB3NzaC1yc2EAAAABIWAAAGEA3FSyQwBI6Z+nCSjUuK8EEAnnkhXlukKoUPND/RRC1Wz2s5TCzIkd3Ou5+Cyz
7   - ssh-rsa AAAAB3NzaC1yc2EAAAABIWAAAQEA3I7VUf215gSn5uavROsc5HRDpZdQueUq5ozemNSj8T7enqKH0EaFoU2VoPgG
8
9 # Send pre-generated ssh private keys to the server
10 # If these are present, they will be written to /etc/ssh and
11 # new random keys will not be generated
12 # in addition to 'rsa' and 'dsa' as shown below, 'ecdsa' is also supported
13 ssh_keys:
14   rsa_private: |
15     -----BEGIN RSA PRIVATE KEY-----
16     MIIBxwIBAAJhAKD0YSHy73nUgysO13XsJmd4fHiFyQ+00R7VVu2iV9Qcon2LZS/x
17     lcydPZ4pQpfjEha6WxZ6o8ci/Ea/wOn+0HGPwaxlEG2Z9inNtj3pgFrYcRztfECb
18     1j6HCibZbAzYtwIBIwJgO8h72WjcmvcpZ8OvHSvTwaGu02TkR6mPgHsgSaKy6GJo
19     PUJnaZRWuba/HX0KGyhz19nPzLpzG5f0fYahlMJAYc13FV7K6kMBPXTRR6FxfgHEg
20     L0MPC7cdqAwOVNcPy6A7AjEA1bNaIjOzFN2sfZX0j7OMhQuc4zP7r80zaGc5oy6W
21     p58hRancFKEvnEq2CeL3vtuZAJEAwNBHpbNsBYTRPCHM7rZuG/iBtwp8Rxhc9I5w
22     ixvzMgi+HpGLWzUIBS+P/XhekIjPAjA285rVmEP+DR255Ls65QbgYhJmTzIXQ2T9
23     luLvcmFBC6l35Uc4gTgg4ALsmXLn7lMCMGMpSWspEvuGInayTCL+vEjmnBT+FAAdO
24     W7D4zCpI43jRS9U06JV0eSc9CDk2lwiA3wIwCTB/6uc8Cq85D9YqpM10FuHjKpnP
25     REPP0yrAspdeOAV+6VKRavstea7+2DZmSUGe
26     -----END RSA PRIVATE KEY-----
27

```

```

28  rsa_public: ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAGEAoPrHIfLvedSDKw7XdewmZ3h8eIXJD7TRHtVW7aJX1ByifYt1L/I
29
30  dsa_private: |
31      -----BEGIN DSA PRIVATE KEY-----
32      MIIBuwIBAAKBgQDP2HLu7pTExL89USyM0264RCyWX/CMLmukxX0Jdbm29ax8FBJT
33      pLrO8TIXVY5rPAJmldTHnpuyJhOvU9G7M8tPUABtzSJh4GVSHlwaCfycwcpLv9TX
34      DgWIpsj+6EiHCyAr1B1/CBp9RiaB+10QcFbm+lapuET+/Au6vSDp9IRt1QIVAIMR
35      8KucvUYbOEI+yv+5LW9u3z/BAoGBAI0q6JP+JvJmwZFaeCMMVxXUbqiSko/P1lsa
36      LNNBHZ5/8MOUIm8rB2FC6ziidfueJpqTMqeQmSA1EBCwnwreUnGfRrKoJpyPNENY
37      dl5MG6N5J+z81sEchFeprryZ+D3Ge9VjPq3Tf3NhKKwCDQ0240aPezbnjPeFm4mH
38      bYxzcZ9GAoGAXmLIFSQgiAPu459rCKxT46tHJtMQQfnNiEnQLbFluefZ/yiI4DI3
39      8UzTCOXLhUA7ybmZha+D/csj15Y9/BNFu07unzVhikCQV9DTeXX46pG4s1o23JKC
40      /QaYWNMZ7kTRv+wWow9MhGiVdML4ZN4Xnifu05krqAybngIy66PMEoQCFEIsKKWv
41      99izizAH0KBMVbxy03Trz
42      -----END DSA PRIVATE KEY-----
43
44  dsa_public: ssh-dss AAAAB3NzaC1kc3MAAACBAM/Ycu7ulMTEvz1RLIzTbrhELJZf8Iwua6TFfQ11ubb1rHwUElOkus7xMh

```

1.6 Datasources

1.6.1 What is a datasource?

Datasources are sources of configuration data for cloud-init that typically come from the user (aka userdata) or come from the stack that created the configuration drive (aka metadata). Typical userdata would include files, yaml, and shell scripts while typical metadata would include server name, instance id, display name and other cloud specific details. Since there are multiple ways to provide this data (each cloud solution seems to prefer its own way) internally a datasource abstract class was created to allow for a single way to access the different cloud systems methods to provide this data through the typical usage of subclasses.

The current interface that a datasource object must provide is the following:

```

# returns a mime multipart message that contains
# all the various fully-expanded components that
# were found from processing the raw userdata string
# - when filtering only the mime messages targeting
#   this instance id will be returned (or messages with
#   no instance id)
def get_userdata(self, apply_filter=False)

# returns the raw userdata string (or none)
def get_userdata_raw(self)

# returns a integer (or none) which can be used to identify
# this instance in a group of instances which are typically
# created from a single command, thus allowing programatic
# filtering on this launch index (or other selective actions)
@property
def launch_index(self)

# the data sources' config_obj is a cloud-config formatted
# object that came to it from ways other than cloud-config
# because cloud-config content would be handled elsewhere
def get_config_obj(self)

#returns a list of public ssh keys

```

```

def get_public_ssh_keys(self)

# translates a device 'short' name into the actual physical device
# fully qualified name (or none if said physical device is not attached
# or does not exist)
def device_name_to_device(self, name)

# gets the locale string this instance should be applying
# which typically used to adjust the instances locale settings files
def get_locale(self)

@property
def availability_zone(self)

# gets the instance id that was assigned to this instance by the
# cloud provider or when said instance id does not exist in the backing
# metadata this will return 'iid-datasource'
def get_instance_id(self)

# gets the fully qualified domain name that this host should be using
# when configuring network or hostname related settings, typically
# assigned either by the cloud provider or the user creating the vm
def get_hostname(self, fqdn=False)

def get_package_mirror_info(self)

```

1.6.2 EC2

The EC2 datasource is the oldest and most widely used datasource that cloud-init supports. This datasource interacts with a *magic* ip that is provided to the instance by the cloud provider. Typically this ip is 169.254.169.254 of which at this ip a http server is provided to the instance so that the instance can make calls to get instance userdata and instance metadata.

Metadata is accessible via the following URL:

```

GET http://169.254.169.254/2009-04-04/meta-data/
ami-id
ami-launch-index
ami-manifest-path
block-device-mapping/
hostname
instance-id
instance-type
local-hostname
local-ipv4
placement/
public-hostname
public-ipv4
public-keys/
reservation-id
security-groups

```

Userdata is accessible via the following URL:

```

GET http://169.254.169.254/2009-04-04/user-data
1234,fred,reboot,true | 4512,jimbo, | 173,,,

```

Note that there are multiple versions of this data provided, cloud-init by default uses **2009-04-04** but newer versions can be supported with relative ease (newer versions have more data exposed, while maintaining backward compatibility with the previous versions).

To see which versions are supported from your cloud provider use the following URL:

```
GET http://169.254.169.254/
1.0
2007-01-19
2007-03-01
2007-08-29
2007-10-10
2007-12-15
2008-02-01
2008-09-01
2009-04-04
...
latest
```

1.6.3 Config Drive

The configuration drive datasource supports the [OpenStack](#) configuration drive disk.

See the [config drive extension](#) and [introduction](#) in the public documentation for more information.

By default, cloud-init does *always* consider this source to be a full-fledged datasource. Instead, the typical behavior is to assume it is really only present to provide networking information. Cloud-init will copy off the network information, apply it to the system, and then continue on. The “full” datasource could then be found in the EC2 metadata service. If this is not the case then the files contained on the located drive must provide equivalents to what the EC2 metadata service would provide (which is typical of the version 2 support listed below)

Version 1

The following criteria are required to as a config drive:

1. Must be formatted with [vfat](#) filesystem
2. Must be a un-partitioned block device (`/dev/vdb`, not `/dev/vdb1`)
3. Must contain *one* of the following files

```
/etc/network/interfaces
/root/.ssh/authorized_keys
/meta.js
```

`/etc/network/interfaces`

This file is laid down by nova in order to pass static networking information to the guest. Cloud-init will copy it off of the config-drive and into `/etc/network/interfaces` (or convert it to RH format) as soon as it can, and then attempt to bring up all network interfaces.

`/root/.ssh/authorized_keys`

This file is laid down by nova, and contains the ssk keys that were provided to nova on instance creation (nova-boot `-key ...`)

`/meta.js`

`meta.js` is populated on the config-drive in response to the user passing “meta flags” (nova boot `-meta key=value ...`). It is expected to be json formatted.

Version 2

The following criteria are required to as a config drive:

1. Must be formatted with `vfat` or `iso9660` filesystem or have a `filesystem` label of **config-2**
2. Must be a un-partitioned block device (`/dev/vdb`, not `/dev/vdb1`)
3. The files that will typically be present in the config drive are:

```

openstack/
- 2012-08-10/ or latest/
  - meta_data.json
  - user_data (not mandatory)
- content/
  - 0000 (referenced content files)
  - 0001
  - ....
ec2
- latest/
  - meta-data.json (not mandatory)

```

Keys and values

Cloud-init's behavior can be modified by keys found in the `meta.js` (version 1 only) file in the following ways.

```

dsmode:
  values: local, net, pass
  default: pass

```

This is what indicates if `configdrive` is a final data source or not. By default it is 'pass', meaning this datasource should not be read. Set it to 'local' or 'net' to stop cloud-init from continuing on to search for other data sources after network config.

The difference between 'local' and 'net' is that local will not require networking to be up before user-data actions (or boothooks) are run.

```

instance-id:
  default: iid-dsconfigdrive

```

This is utilized as the metadata's instance-id. It should generally be unique, as it is what is used to determine "is this a new instance".

```

public-keys:
  default: None

```

If present, these keys will be used as the public keys for the instance. This value overrides the content in `authorized_keys`.

Note: it is likely preferable to provide keys via user-data

```

user-data:
  default: None

```

This provides cloud-init user-data. See *examples* for what all can be present here.

1.6.4 OpenNebula

The [OpenNebula](#) (ON) datasource supports the contextualization disk.

See [contextualization overview](#), [contextualizing VMs](#) and [network configuration](#) in the public documentation for more information.

OpenNebula's virtual machines are contextualized (parametrized) by CD-ROM image, which contains a shell script *context.sh* with custom variables defined on virtual machine start. There are no fixed contextualization variables, but the datasource accepts many used and recommended across the documentation.

Datasource configuration

Datasource accepts following configuration options.

```
dsmode:
  values: local, net, disabled
  default: net
```

Tells if this datasource will be processed in 'local' (pre-networking) or 'net' (post-networking) stage or even completely 'disabled'.

```
parseuser:
  default: nobody
```

Unprivileged system user used for contextualization script processing.

Contextualization disk

The following criteria are required:

1. Must be formatted with *iso9660* filesystem or have a *filesystem* label of **CONTEXT** or **CDROM**
2. Must contain file *context.sh* with contextualization variables. File is generated by OpenNebula, it has a `KEY='VALUE'` format and can be easily read by bash

Contextualization variables

There are no fixed contextualization variables in OpenNebula, no standard. Following variables were found on various places and revisions of the OpenNebula documentation. Where multiple similar variables are specified, only first found is taken.

```
DSMODE
```

Datasource mode configuration override. Values: local, net, disabled.

```
DNS
ETH<x>_IP
ETH<x>_NETWORK
ETH<x>_MASK
ETH<x>_GATEWAY
ETH<x>_DOMAIN
ETH<x>_DNS
```

Static [network configuration](#).

```
HOSTNAME
```

Instance hostname.

```
PUBLIC_IP
IP_PUBLIC
ETH0_IP
```

If no hostname has been specified, cloud-init will try to create hostname from instance's IP address in 'local' dsmode. In 'net' dsmode, cloud-init tries to resolve one of its IP addresses to get hostname.

```
SSH_KEY
SSH_PUBLIC_KEY
```

One or multiple SSH keys (separated by newlines) can be specified.

```
USER_DATA
USERDATA
```

cloud-init user data.

Example configuration

This example cloud-init configuration (*cloud.cfg*) enables OpenNebula datasource only in 'net' mode.

```
disable_ec2_metadata: True
datasource_list: ['OpenNebula']
datasource:
  OpenNebula:
    dsmode: net
    parseuser: nobody
```

Example VM's context section

```
CONTEXT=[
  PUBLIC_IP="$NIC[IP] ",
  SSH_KEY="$USER[SSH_KEY]
$USER[SSH_KEY1]
$USER[SSH_KEY2] ",
  USER_DATA="#cloud-config
# see https://help.ubuntu.com/community/CloudInit

packages: []

mounts:
- [vdc, none, swap, sw, 0, 0]
runcmd:
- echo 'Instance has been configured by cloud-init.' | wall
" ]
```

1.6.5 Alt cloud

The datasource altcloud will be used to pick up user data on RHEVm and vSphere.

RHEVm

For RHEVm v3.0 the userdata is injected into the VM using floppy injection via the RHEVm dashboard "Custom Properties".

The format of the Custom Properties entry must be:

```
floppyinject=user-data.txt:<base64 encoded data>
```

For example to pass a simple bash script:

```
% cat simple_script.bash
#!/bin/bash
echo "Hello Joe!" >> /tmp/JJV_Joe_out.txt

% base64 < simple_script.bash
IyEvYmluL2Jhc2gKZWNoYAiSGVsbG8gSm9lISIgPj4gL3RtcC9KS1ZfSm91X291dC50eHQK
```

To pass this example script to cloud-init running in a RHEVm v3.0 VM set the “Custom Properties” when creating the RHEMv v3.0 VM to:

```
floppyinject=user-data.txt:IyEvYmluL2Jhc2gKZWNoYAiSGVsbG8gSm9lISIgPj4gL3RtcC9KS1ZfSm91X291dC50eHQK
```

NOTE: The prefix with file name must be: floppyinject=user-data.txt:

It is also possible to launch a RHEVm v3.0 VM and pass optional user data to it using the Delta Cloud.

For more information on Delta Cloud see: <http://deltacloud.apache.org>

vSphere

For VMWare’s vSphere the userdata is injected into the VM as an ISO via the cdrom. This can be done using the vSphere dashboard by connecting an ISO image to the CD/DVD drive.

To pass this example script to cloud-init running in a vSphere VM set the CD/DVD drive when creating the vSphere VM to point to an ISO on the data store.

Note: The ISO must contain the user data.

For example, to pass the same simple_script.bash to vSphere:

Create the ISO

```
% mkdir my-iso
```

NOTE: The file name on the ISO must be: user-data.txt

```
% cp simple_scrip.t.bash my-iso/user-data.txt
% genisoimage -o user-data.iso -r my-iso
```

Verify the ISO

```
% sudo mkdir /media/vsphere_iso
% sudo mount -o loop JoeV_CI_02.iso /media/vsphere_iso
% cat /media/vsphere_iso/user-data.txt
% sudo umount /media/vsphere_iso
```

Then, launch the vSphere VM the ISO user-data.iso attached as a CDRROM.

It is also possible to launch a vSphere VM and pass optional user data to it using the Delta Cloud.

For more information on Delta Cloud see: <http://deltacloud.apache.org>

1.6.6 No cloud

The data source `NoCloud` and `NoCloudNet` allow the user to provide user-data and meta-data to the instance without running a network service (or even without having a network at all).

You can provide meta-data and user-data to a local vm boot via files on a `vfat` or `iso9660` filesystem. The filesystem volume label must be `cidata`.

These user-data and meta-data files are expected to be in the following format.

```
/user-data
/meta-data
```

Basically, user-data is simply user-data and meta-data is a yaml formatted file representing what you'd find in the EC2 metadata service.

Given a disk `ubuntu 12.04 cloud image` in `'disk.img'`, you can create a sufficient disk by following the example below.

```
## create user-data and meta-data files that will be used
## to modify image on first boot
$ { echo instance-id: iid-local01; echo local-hostname: cloudimg; } > meta-data

$ printf "#cloud-config\npassword: passw0rd\nchpasswd: { expire: False }\nssh_pwauth: True\n" > user-

## create a disk to attach with some user-data and meta-data
$ genisoimage -output seed.iso -volid cidata -joliet -rock user-data meta-data

## alternatively, create a vfat filesystem with same files
## $ truncate --size 2M seed.img
## $ mkfs.vfat -n cidata seed.img
## $ mcopy -oi seed.img user-data meta-data ::

## create a new qcow image to boot, backed by your original image
$ qemu-img create -f qcow2 -b disk.img boot-disk.img

## boot the image and login as 'ubuntu' with password 'passw0rd'
## note, passw0rd was set as password through the user-data above,
## there is no password set on these images.
$ kvm -m 256 \
  -net nic -net user,hostfwd=tcp::2222-:22 \
  -drive file=boot-disk.img,if=virtio \
  -drive file=seed.iso,if=virtio
```

Note: that the instance-id provided (`iid-local01` above) is what is used to determine if this is “first boot”. So if you are making updates to user-data you will also have to change that, or start the disk fresh.

Also, you can inject an `/etc/network/interfaces` file by providing the content for that file in the `network-interfaces` field of metadata.

Example metadata:

```
instance-id: iid-abcdefg
network-interfaces: |
  iface eth0 inet static
  address 192.168.1.10
  network 192.168.1.0
  netmask 255.255.255.0
  broadcast 192.168.1.255
  gateway 192.168.1.254
hostname: myhost
```

1.6.7 MAAS

TODO

For now see: <http://maas.ubuntu.com/>

1.6.8 CloudStack

Apache CloudStack expose user-data, meta-data, user password and account sshkey thru the Virtual-Router. For more details on meta-data and user-data, refer the [CloudStack Administrator Guide](#).

URLs to access user-data and meta-data from the Virtual Machine. Here 10.1.1.1 is the Virtual Router IP:

```
http://10.1.1.1/latest/user-data
http://10.1.1.1/latest/meta-data
http://10.1.1.1/latest/meta-data/{metadata type}
```

Configuration

Apache CloudStack datasource can be configured as follows:

```
datasource:
  CloudStack: {}
  None: {}
datasource_list:
  - CloudStack
```

1.6.9 OVF

TODO

For now see: <https://bazaar.launchpad.net/~cloud-init-dev/cloud-init/trunk/files/head:/doc/sources/ovf/>

1.6.10 OpenStack

TODO

Vendor Data

The OpenStack metadata server can be configured to serve up vendor data which is available to all instances for consumption. OpenStack vendor data is, generally, a JSON object.

cloud-init will look for configuration in the `cloud-init` attribute of the vendor data JSON object. cloud-init processes this configuration using the same handlers as user data, so any formats that work for user data should work for vendor data.

For example, configuring the following as vendor data in OpenStack would upgrade packages and install `htop` on all instances:

```
{"cloud-init": "#cloud-config\npackage_upgrade: True\npackages:\n - htop"}
```

For more general information about how cloud-init handles vendor data, including how it can be disabled by users on instances, see <https://bazaar.launchpad.net/~cloud-init-dev/cloud-init/trunk/view/head:/doc/vendordata.txt>

1.6.11 Fallback/None

This is the fallback datasource when no other datasource can be selected. It is the equivalent of a *empty* datasource in that it provides a empty string as userdata and a empty dictionary as metadata. It is useful for testing as well as for when you do not have a need to have an actual datasource to meet your instance requirements (ie you just want to run modules that are not concerned with any external data). It is typically put at the end of the datasource search list so that if all other datasources are not matched, then this one will be so that the user is not left with an inaccessible instance.

Note: the instance id that this datasource provides is `iid-datasource-none`.

1.7 Modules

1.7.1 Apt Configure

Internal name: `cc_apt_configure`

1.7.2 Apt Pipelining

Internal name: `cc_apt_pipelining`

1.7.3 Bootcmd

Internal name: `cc_bootcmd`

1.7.4 Byobu

Internal name: `cc_byobu`

1.7.5 Ca Certs

Internal name: `cc_ca_certs`

1.7.6 Chef

Internal name: `cc_chef` **Summary:** module that configures, starts and installs chef.

Description: This module enables chef to be installed (from packages or from gems, or from omnibus). Before this occurs chef configurations are written to disk (`validation.pem`, `client.pem`, `firstboot.json`, `client.rb`), and needed chef folders/directories are created (`/etc/chef` and `/var/log/chef` and so-on). Then once installing proceeds correctly if configured chef will be started (in daemon mode or in non-daemon mode) and then once that has finished (if ran in non-daemon mode this will be when chef finishes converging, if ran in daemon mode then no further actions are possible since chef will have forked into its own process) then a post run function can run that can do finishing activities (such as removing the validation pem file).

It can be configured with the following option structure:

```
chef:
  directories: (defaulting to /etc/chef, /var/log/chef, /var/lib/chef,
               /var/cache/chef, /var/backups/chef, /var/run/chef)
  validation_cert: (optional string to be written to file validation_key)
                   special value 'system' means set use existing file
  validation_key: (optional the path for validation_cert. default
                 /etc/chef/validation.pem)
  firstboot_path: (path to write run_list and initial_attributes keys that
                 should also be present in this configuration, defaults
                 to /etc/chef/firstboot.json)
  exec: boolean to run or not run chef (defaults to false, unless
      a gem installed is requested
      where this will then default
      to true)

chef.rb template keys (if falsey, then will be skipped and not
                       written to /etc/chef/client.rb)
```

```
chef:
  client_key:
  environment:
  file_backup_path:
  file_cache_path:
  json_attribs:
  log_level:
  log_location:
  node_name:
  pid_file:
  server_url:
  show_time:
  ssl_verify_mode:
  validation_cert:
  validation_key:
  validation_name:
```

`cloudinit.config.cc_chef.handle` (*name, cfg, cloud, log, _args*)
Handler method activated by cloud-init.

1.7.7 Debug

Internal name: `cc_debug` **Summary:** helper to debug cloud-init *internal* datastructures.

Description: This module will enable for outputting various internal information that cloud-init sources provide to either a file or to the output console/log location that this cloud-init has been configured with when running.

It can be configured with the following option structure:

```
debug:
  verbose: (defaulting to true)
  output: (location to write output, defaulting to console + log)
```

Note: Log configurations are not output.

`cloudinit.config.cc_debug.handle` (*name, cfg, cloud, log, args*)
Handler method activated by cloud-init.

1.7.8 Disable Ec2 Metadata

Internal name: cc_disable_ec2_metadata

1.7.9 Disk Setup

Internal name: cc_disk_setup

1.7.10 Emit Upstart

Internal name: cc_emit_upstart

1.7.11 Final Message

Internal name: cc_final_message

1.7.12 Foo

Internal name: cc_foo

1.7.13 Growpart

Internal name: cc_growpart

1.7.14 Grub Dpkg

Internal name: cc_grub_dpkg

1.7.15 Keys To Console

Internal name: cc_keys_to_console

1.7.16 Landscape

Internal name: cc_landscape

1.7.17 Locale

Internal name: cc_locale

1.7.18 Mcollective

Internal name: cc_mcollective

1.7.19 Migrator

Internal name: `cc_migrator`

1.7.20 Mounts

Internal name: `cc_mounts`

1.7.21 Package Update Upgrade Install

Internal name: `cc_package_update_upgrade_install`

1.7.22 Phone Home

Internal name: `cc_phone_home`

1.7.23 Power State Change

Internal name: `cc_power_state_change`

1.7.24 Puppet

Internal name: `cc_puppet`

1.7.25 Resizesfs

Internal name: `cc_resizesfs`

1.7.26 Resolv Conf

Internal name: `cc_resolv_conf`

1.7.27 Rightscale Userdata

Internal name: `cc_rightscale_userdata`

1.7.28 Rsyslog

Internal name: `cc_rsyslog` rsyslog module allows configuration of syslog logging via rsyslog Configuration is done under the cloud-config top level 'rsyslog'.

Under 'rsyslog' you can define:

- `configs`: [default=[]] this is a list. entries in it are a string or a dictionary. each entry has 2 parts:
 - `content`
 - `filename`

if the entry is a string, then it is assigned to 'content'. for each entry, content is written to the provided filename. if filename is not provided, its default is read from 'config_filename'

Content here can be any valid rsyslog configuration. No format specific format is enforced.

For simply logging to an existing remote syslog server, via udp: configs: [". @192.168.1.1"]

- remotes: [default={}] This is a dictionary of name / value pairs. In comparison to 'config's, it is more focused in that it only supports remote syslog configuration. It is not rsyslog specific, and could convert to other syslog implementations.

Each entry in remotes is a 'name' and a 'value'.

- name: an string identifying the entry. good practice would indicate using a consistent and identifiable string for the producer. For example, the MAAS service could use 'maas' as the key.
- value consists of the following parts:
 - * optional filter for log messages
 - default if not present: .
 - * optional leading '@' or '@@' (indicates udp or tcp respectively). default if not present (udp): @ This is rsyslog format for that. if not present, is '@'.
 - * ipv4 or ipv6 or hostname ipv6 addresses must be in [::1] format. (@[fd00::1]:514)
 - * optional port defaults to 514

- config_filename: [default=20-cloud-config.conf] this is the file name to use if none is provided in a config entry.
- config_dir: [default=/etc/rsyslog.d] this directory is used for filenames that are not absolute paths.
- service_reload_command: [default="auto"] this command is executed if files have been written and thus the syslog daemon needs to be told.

Note, since cloud-init 0.5 a legacy version of rsyslog config has been present and is still supported. See below for the mappings between old value and new value:

```
old value -> new value
'rsyslog' -> rsyslog/configs 'rsyslog_filename' -> rsyslog/config_filename
'rsyslog_dir' -> rsyslog/config_dir
```

the legacy config does not support 'service_reload_command'.

Example config: #cloud-config rsyslog:

configs:

- “. @@192.158.1.1”
- content: “. @@192.0.2.1:10514” filename: 01-example.conf
- content: |. @syslogd.example.com

remotes: maas: “192.168.1.1” juju: “10.0.4.1”

config_dir: config_dir config_filename: config_filename service_reload_command: [your, syslog, restart, command]

Example Legacy config: #cloud-config rsyslog:

- “. @@192.158.1.1”

rsyslog_dir: /etc/rsyslog-config.d/ rsyslog_filename: 99-local.conf

1.7.29 Runcmd

Internal name: `cc_runcmd`

1.7.30 Salt Minion

Internal name: `cc_salt_minion`

1.7.31 Scripts Per Boot

Internal name: `cc_scripts_per_boot`

1.7.32 Scripts Per Instance

Internal name: `cc_scripts_per_instance`

1.7.33 Scripts Per Once

Internal name: `cc_scripts_per_once`

1.7.34 Scripts User

Internal name: `cc_scripts_user`

1.7.35 Scripts Vendor

Internal name: `cc_scripts_vendor`

1.7.36 Seed Random

Internal name: `cc_seed_random`

1.7.37 Set Hostname

Internal name: `cc_set_hostname`

1.7.38 Set Passwords

Internal name: `cc_set_passwords`

1.7.39 Ssh

Internal name: `cc_ssh`

1.7.40 Ssh Authkey Fingerprints

Internal name: `cc_ssh_authkey_fingerprints`

1.7.41 Ssh Import Id

Internal name: `cc_ssh_import_id`

1.7.42 Timezone

Internal name: `cc_timezone`

1.7.43 Ubuntu Init Switch

Internal name: `cc_ubuntu_init_switch` **Summary:** reboot system into another init.

Description: This module provides a way for the user to boot with systemd even if the image is set to boot with upstart. It should be run as one of the first `cloud_init_modules`, and will switch the init system and then issue a reboot. The next boot will come up in the target init system and no action will be taken.

This should be inert on non-ubuntu systems, and also exit quickly.

It can be configured with the following option structure:

```
init_switch:
  target: systemd (can be 'systemd' or 'upstart')
  reboot: true (reboot if a change was made, or false to not reboot)
```

Note: Best effort is made, but it's possible this system will break, and probably won't interact well with any other mechanism you've used to switch the init system.

`cloudinit.config.cc_ubuntu_init_switch.handle` (*name, cfg, cloud, log, args*)
 Handler method activated by cloud-init.

1.7.44 Update Etc Hosts

Internal name: `cc_update_etc_hosts`

1.7.45 Update Hostname

Internal name: `cc_update_hostname`

1.7.46 Users Groups

Internal name: `cc_users_groups`

1.7.47 Write Files

Internal name: `cc_write_files`

1.7.48 Yum Add Repo

Internal name: cc_yum_add_repo

1.8 Merging User-Data Sections

1.8.1 Overview

This was implemented because it has been a common feature request that there be a way to specify how cloud-config yaml “dictionaries” provided as user-data are merged together when there are multiple yamls to merge together (say when performing an #include).

Since previously the merging algorithm was very simple and would only overwrite and not append lists, or strings, and so on it was decided to create a new and improved way to merge dictionaries (and there contained objects) together in a way that is customizable, thus allowing for users who provide cloud-config user-data to determine exactly how there objects will be merged.

For example.

```
#cloud-config (1)
run_cmd:
  - bash1
  - bash2

#cloud-config (2)
run_cmd:
  - bash3
  - bash4
```

The previous way of merging the following 2 objects would result in a final cloud-config object that contains the following.

```
#cloud-config (merged)
run_cmd:
  - bash3
  - bash4
```

Typically this is not what users want, instead they would likely prefer:

```
#cloud-config (merged)
run_cmd:
  - bash1
  - bash2
  - bash3
  - bash4
```

This way makes it easier to combine the various cloud-config objects you have into a more useful list, thus reducing duplication that would have had to occur in the previous method to accomplish the same result.

1.8.2 Customizability

Since the above merging algorithm may not always be the desired merging algorithm (like how the previous merging algorithm was not always the preferred one) the concept of customizing how merging can be done was introduced through a new concept call ‘merge classes’.

A merge class is a class definition which provides functions that can be used to merge a given type with another given type.

An example of one of these merging classes is the following:

```
class Merger(object):
    def __init__(self, merger, opts):
        self._merger = merger
        self._overwrite = 'overwrite' in opts

    # This merging algorithm will attempt to merge with
    # another dictionary, on encountering any other type of object
    # it will not merge with said object, but will instead return
    # the original value
    #
    # On encountering a dictionary, it will create a new dictionary
    # composed of the original and the one to merge with, if 'overwrite'
    # is enabled then keys that exist in the original will be overwritten
    # by keys in the one to merge with (and associated values). Otherwise
    # if not in overwrite mode the 2 conflicting keys themselves will
    # be merged.
    def _on_dict(self, value, merge_with):
        if not isinstance(merge_with, (dict)):
            return value
        merged = dict(value)
        for (k, v) in merge_with.items():
            if k in merged:
                if not self._overwrite:
                    merged[k] = self._merger.merge(merged[k], v)
                else:
                    merged[k] = v
            else:
                merged[k] = v
        return merged
```

As you can see there is a ‘_on_dict’ method here that will be given a source value and a value to merge with. The result will be the merged object. This code itself is called by another merging class which ‘directs’ the merging to happen by analyzing the types of the objects to merge and attempting to find a known object that will merge that type. I will avoid pasting that here, but it can be found in the *mergers/__init__.py* file (see *LookupMerger* and *UnknownMerger*).

So following the typical cloud-init way of allowing source code to be downloaded and used dynamically, it is possible for users to inject their own merging files to handle specific types of merging as they choose (the basic ones included will handle lists, dicts, and strings). Note how each merge can have options associated with it which affect how the merging is performed, for example a dictionary merger can be told to overwrite instead of attempt to merge, or a string merger can be told to append strings instead of discarding other strings to merge with.

1.8.3 How to activate

There are a few ways to activate the merging algorithms, and to customize them for your own usage.

1. The first way involves the usage of MIME messages in cloud-init to specify multipart documents (this is one way in which multiple cloud-config is joined together into a single cloud-config). Two new headers are looked for, both of which can define the way merging is done (the first header to exist wins). These new headers (in lookup order) are ‘Merge-Type’ and ‘X-Merge-Type’. The value should be a string which will satisfy the new merging format definition (see below for this format).
2. The second way is actually specifying the merge-type in the body of the cloud-config dictionary. There are 2 ways to specify this, either as a string or as a dictionary (see format below). The keys that are looked up for this

definition are the following (in order), 'merge_how', 'merge_type'.

String format

The string format that is expected is the following.

```
classname1(option1,option2)+classname2(option3,option4)....
```

The class name there will be connected to class names used when looking for the class that can be used to merge and options provided will be given to the class on construction of that class.

For example, the default string that is used when none is provided is the following:

```
list()+dict()+str()
```

Dictionary format

In cases where a dictionary can be used to specify the same information as the string format (ie option #2 of above) it can be used, for example.

```
{'merge_how': [{'name': 'list', 'settings': ['extend']},  
               {'name': 'dict', 'settings': []},  
               {'name': 'str', 'settings': ['append']}]}
```

This would be the equivalent format for default string format but in dictionary form instead of string form.

1.8.4 Specifying multiple types and its effect

Now you may be asking yourself, if I specify a merge-type header or dictionary for every cloud-config that I provide, what exactly happens?

The answer is that when merging, a stack of 'merging classes' is kept, the first one on that stack is the default merging classes, this set of mergers will be used when the first cloud-config is merged with the initial empty cloud-config dictionary. If the cloud-config that was just merged provided a set of merging classes (via the above formats) then those merging classes will be pushed onto the stack. Now if there is a second cloud-config to be merged then the merging classes from the cloud-config before the first will be used (not the default) and so on. This way a cloud-config can decide how it will merge with a cloud-config dictionary coming after it.

1.8.5 Other uses

In addition to being used for merging user-data sections, the default merging algorithm for merging 'conf.d' yaml files (which form an initial yaml config for cloud-init) was also changed to use this mechanism so its full benefits (and customization) can also be used there as well. Other places that used the previous merging are also, similarly, now extensible (metadata merging, for example).

Note, however, that merge algorithms are not used *across* types of configuration. As was the case before merging was implemented, user-data will overwrite conf.d configuration without merging.

1.9 More information

1.9.1 Useful external references

- [The beauty of cloudinit](#)
- [Introduction to cloud-init \(video\)](#)

1.10 Hacking on cloud-init

This document describes how to contribute changes to cloud-init.

1.10.1 Do these things once

- If you have not already, be sure to sign the CCA:
 - [Canonical Contributor Agreement](#)

- Clone the LaunchPad repository:

```
git clone YOUR_USERNAME@git.launchpad.net:cloud-init cd cloud-init
```

If you would prefer a bazaar style *git clone lp:cloud-init*, see the [Instructions on LaunchPad](#) for more information.

- Create a new remote pointing to your personal LaunchPad repository:

```
git remote add YOUR_USERNAME YOUR_USERNAME@git.launchpad.net:~YOUR_USERNAME/cloud-init
```

1.10.2 Do these things for each feature or bug

- Create a new topic branch for your work:

```
git checkout -b my-topic-branch
```

- Make and commit your changes (note, you can make multiple commits, fixes, more commits.):

```
git commit
```

- Check pep8 and test, and address any issues:

```
make test pep8
```

- Push your changes to your personal LaunchPad repository:

```
git push -u YOUR_USERNAME my-topic-branch
```

- Use your browser to create a merge request:

- Open the branch on [LaunchPad](#)

* It will typically be at https://code.launchpad.net/~YOUR_USERNAME/cloud-init/+git/cloud-init for example <https://code.launchpad.net/~larsks/cloud-init/+git/cloud-init/+ref/feature/move-to-git>

- Click ‘Propose for merging’
- Select `cloud-init` as the target repository
- Select `master` as the target reference path

Then, someone on cloud-init-dev (currently [Scott Moser](#) and [Joshua Harlow](#)) will review your changes and follow up in the merge request.

Feel free to ping and/or join `#cloud-init` on freenode (irc) if you have any questions.

C

cloudinit.config.cc_apt_configure, 29
cloudinit.config.cc_apt_pipelining, 29
cloudinit.config.cc_bootcmd, 29
cloudinit.config.cc_byobu, 29
cloudinit.config.cc_ca_certs, 29
cloudinit.config.cc_chef, 29
cloudinit.config.cc_debug, 30
cloudinit.config.cc_disable_ec2_metadata,
31
cloudinit.config.cc_disk_setup, 31
cloudinit.config.cc_emit_upstart, 31
cloudinit.config.cc_final_message, 31
cloudinit.config.cc_foo, 31
cloudinit.config.cc_growpart, 31
cloudinit.config.cc_grub_dpkg, 31
cloudinit.config.cc_keys_to_console, 31
cloudinit.config.cc_landscape, 31
cloudinit.config.cc_locale, 31
cloudinit.config.cc_mcollective, 31
cloudinit.config.cc_migrator, 32
cloudinit.config.cc_mounts, 32
cloudinit.config.cc_package_update_upgrade_install,
32
cloudinit.config.cc_phone_home, 32
cloudinit.config.cc_power_state_change,
32
cloudinit.config.cc_puppet, 32
cloudinit.config.cc_resizefs, 32
cloudinit.config.cc_resolv_conf, 32
cloudinit.config.cc_rightscale_userdata,
32
cloudinit.config.cc_rsyslog, 32
cloudinit.config.cc_runcmd, 34
cloudinit.config.cc_salt_minion, 34
cloudinit.config.cc_scripts_per_boot,
34
cloudinit.config.cc_scripts_per_instance,
34
cloudinit.config.cc_scripts_per_once,
34
cloudinit.config.cc_scripts_user, 34
cloudinit.config.cc_scripts_vendor, 34
cloudinit.config.cc_seed_random, 34
cloudinit.config.cc_set_hostname, 34
cloudinit.config.cc_set_passwords, 34
cloudinit.config.cc_ssh, 34
cloudinit.config.cc_ssh_authkey_fingerprints,
35
cloudinit.config.cc_ssh_import_id, 35
cloudinit.config.cc_timezone, 35
cloudinit.config.cc_ubuntu_init_switch,
35
cloudinit.config.cc_update_etc_hosts,
35
cloudinit.config.cc_update_hostname, 35
cloudinit.config.cc_users_groups, 35
cloudinit.config.cc_write_files, 35
cloudinit.config.cc_yum_add_repo, 36

C

cloudinit.config.cc_apt_configure (module), 29
 cloudinit.config.cc_apt_pipelining (module), 29
 cloudinit.config.cc_bootcmd (module), 29
 cloudinit.config.cc_byobu (module), 29
 cloudinit.config.cc_ca_certs (module), 29
 cloudinit.config.cc_chef (module), 29
 cloudinit.config.cc_debug (module), 30
 cloudinit.config.cc_disable_ec2_metadata (module), 31
 cloudinit.config.cc_disk_setup (module), 31
 cloudinit.config.cc_emit_upstart (module), 31
 cloudinit.config.cc_final_message (module), 31
 cloudinit.config.cc_foo (module), 31
 cloudinit.config.cc_growpart (module), 31
 cloudinit.config.cc_grub_dpkg (module), 31
 cloudinit.config.cc_keys_to_console (module), 31
 cloudinit.config.cc_landscape (module), 31
 cloudinit.config.cc_locale (module), 31
 cloudinit.config.cc_mcollective (module), 31
 cloudinit.config.cc_migrator (module), 32
 cloudinit.config.cc_mounts (module), 32
 cloudinit.config.cc_package_update_upgrade_install
 (module), 32
 cloudinit.config.cc_phone_home (module), 32
 cloudinit.config.cc_power_state_change (module), 32
 cloudinit.config.cc_puppet (module), 32
 cloudinit.config.cc_resizefs (module), 32
 cloudinit.config.cc_resolv_conf (module), 32
 cloudinit.config.cc_rightscale_userdata (module), 32
 cloudinit.config.cc_rsyslog (module), 32
 cloudinit.config.cc_runcmd (module), 34
 cloudinit.config.cc_salt_minion (module), 34
 cloudinit.config.cc_scripts_per_boot (module), 34
 cloudinit.config.cc_scripts_per_instance (module), 34
 cloudinit.config.cc_scripts_per_once (module), 34
 cloudinit.config.cc_scripts_user (module), 34
 cloudinit.config.cc_scripts_vendor (module), 34
 cloudinit.config.cc_seed_random (module), 34
 cloudinit.config.cc_set_hostname (module), 34
 cloudinit.config.cc_set_passwords (module), 34

cloudinit.config.cc_ssh (module), 34
 cloudinit.config.cc_ssh_authkey_fingerprints (module),
 35
 cloudinit.config.cc_ssh_import_id (module), 35
 cloudinit.config.cc_timezone (module), 35
 cloudinit.config.cc_ubuntu_init_switch (module), 35
 cloudinit.config.cc_update_etc_hosts (module), 35
 cloudinit.config.cc_update_hostname (module), 35
 cloudinit.config.cc_users_groups (module), 35
 cloudinit.config.cc_write_files (module), 35
 cloudinit.config.cc_yum_add_repo (module), 36

H

handle() (in module cloudinit.config.cc_chef), 30
 handle() (in module cloudinit.config.cc_debug), 30
 handle() (in module cloudinit.config.cc_ubuntu_init_switch), 35