ubuntu

Building physical clouds

Christian "kiko" Reis
David Duffey
Luke Williams

OCP Tech Day
August 30th, Menlo Park, CA

CANONICAL

Preamble

How VMware killed hardware

Time from request to prompt

VM

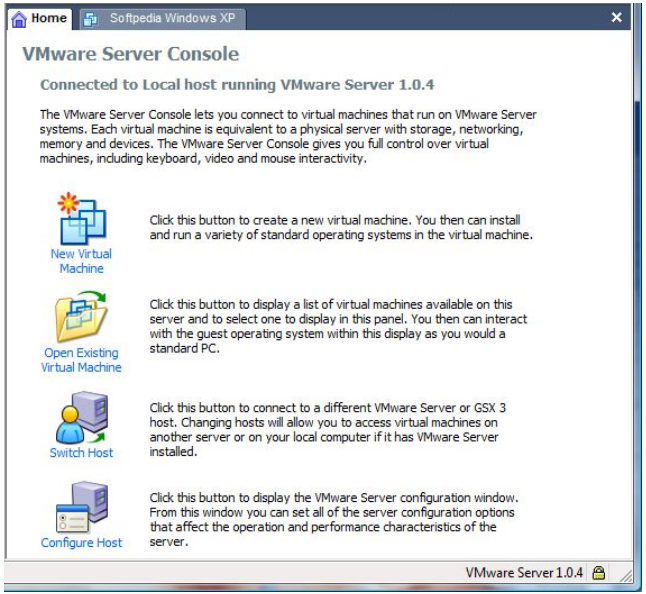

Metal

Request
Procure
Rack
Firmware
Configure
Install OS
Launch
Credentials
Automation

VM

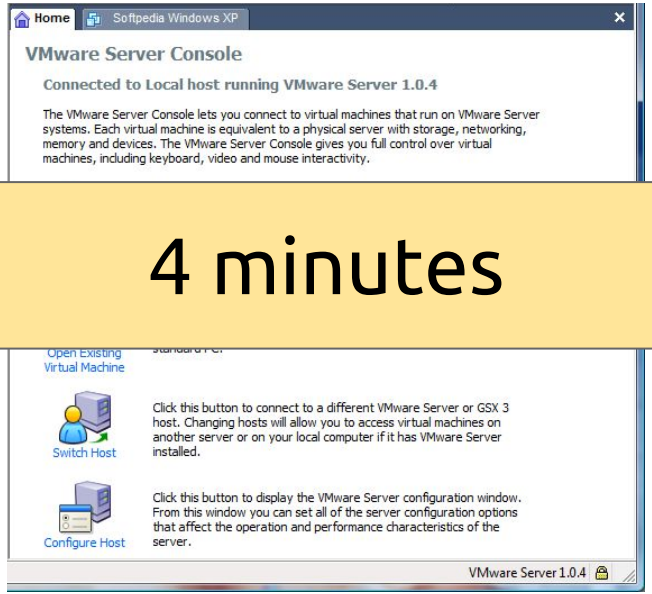

4 minutes

Metal

Request
Procure
Rack
Install OS
Launch
Credentials
Automation

4 months

Cloud

```
>>> from boto import ec2
>>> c = ec2.connect('us-west-2')
>>> r = c.run_instances(
     min_count=3,
     aws_image_id,
     instance_type='c1.xlarge')
```

Metal

Request
Procure
Rack
Firmware
Configure
Install OS
Launch
Credentials
Automation

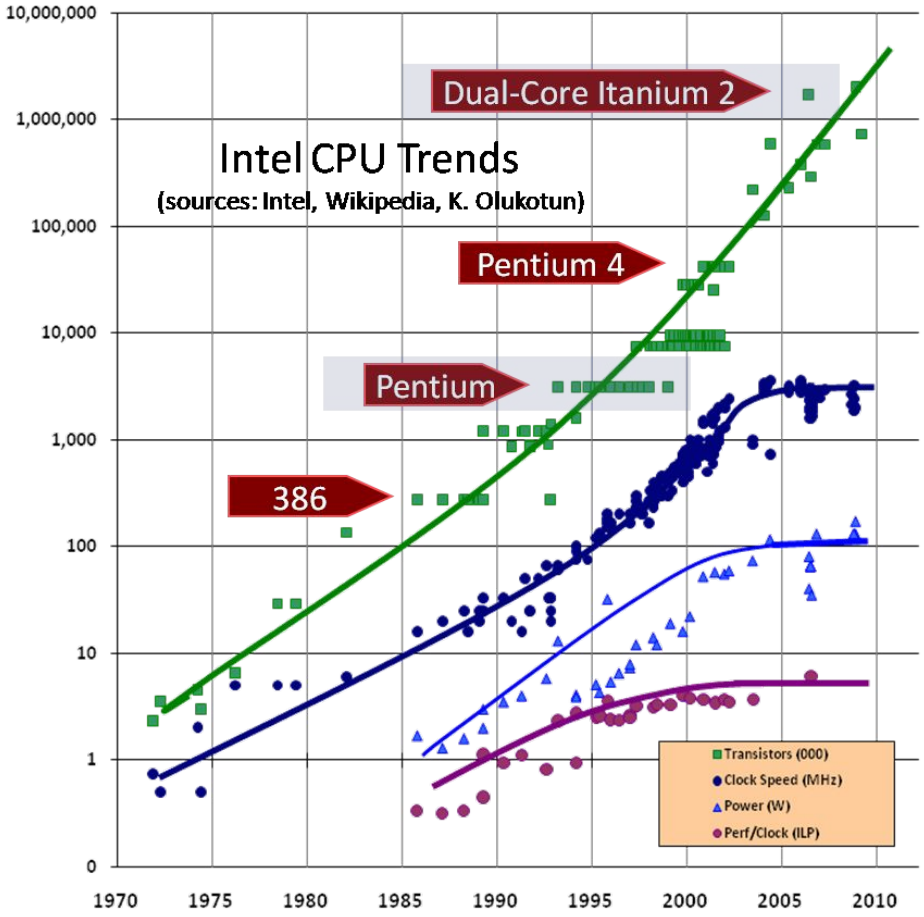
Intel CPU Trends
(sources: Intel, Wikipedia, K. Olukotun)
Dual-Core Itanium 2
Pentium 4
Pentium
386
Transistors (000)
Clock Speed (MHz)
Power (W)
Perf/Clock (ILP)
10,000,000
1,000,000
100,000
10,000
1,000
100
10
1
0
1970
1975
1980
1985
1990
1995
2000
2005
2010

Big Software (noun)

a complex assembly of many software components sourced from different vendors, running on multiple machines, providing to its user the impression of a single system

OCP

1. Focus on versatile, proven system designs
2. Pooling of talent to ensure innovation
3. Applying YAGNI to monitoring & firmware
4. Emphasizing operational efficiency

OCP

1. Focus on versatile, proven system designs
2. Pooling of talent to ensure innovation
3. Applying YAGNI to monitoring & firmware
4. Emphasizing operational efficiency

MAAS

An open source API
for hardware operations

MAAS

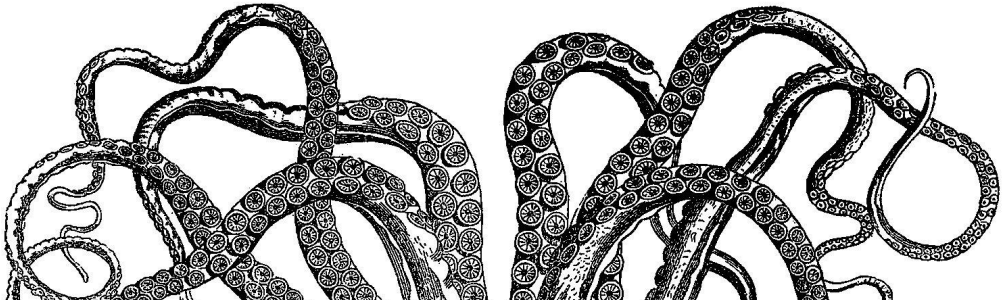

MAAS

PXE DHCP tftp iSCSI IPMI AMT iLO DRAC UCSM Redfish

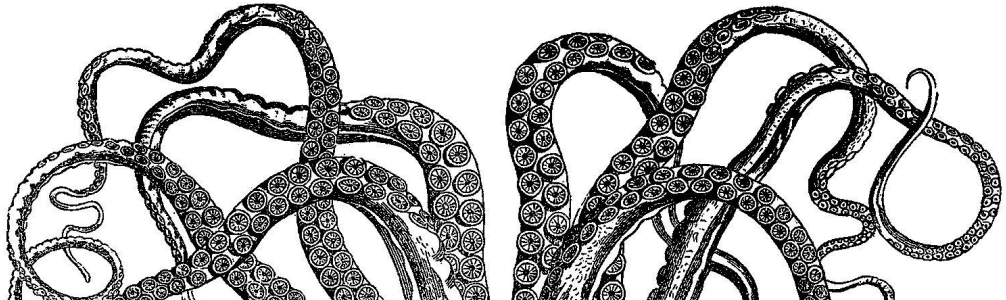

DNS HTTP REST SSH

MAAS

PXE DHCP tftp iSCSI IPMI AMT iLO DRAC UCSM Redfish

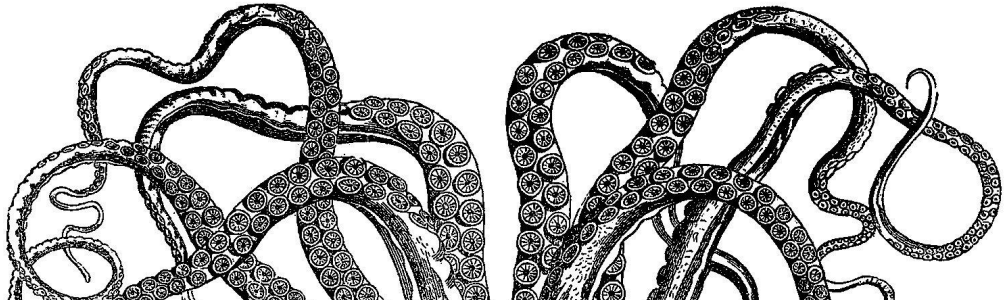





DNS HTTP REST SSH

MAAS

PXE DHCP tftp iSCSI IPMI AMT iLO DRAC UCSM Redfish

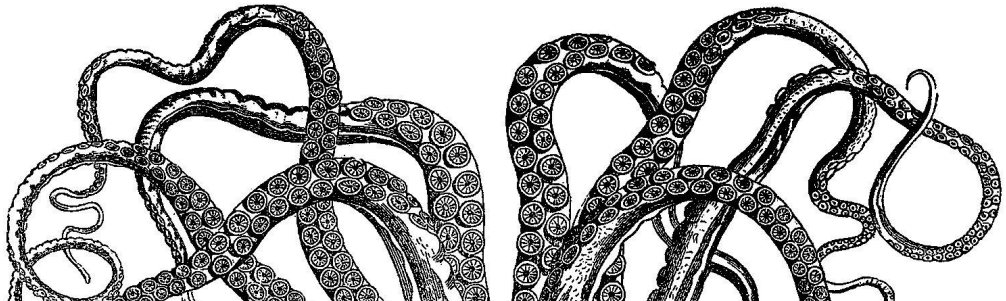

```
$ maas login devcluster http://maas.dev/MAAS/api/2.0

$ maas devcluster node start distro_series="centos7"
```

facebook MAAS

5 Nodes | 0 Devices

Add Hardware

Filter by

- Status
- Owner
- Tags
  - autopilot (1)
  - landscape (1)
  - node1 (1)
  - node2 (1)
  - node3 (1)
  - node4 (1)
  - node5 (1)
- Storage Tags
- Subnets
- Fabrics

Search nodes

FQDN | MAC	Power	Status	Owner	Cores	RAM (GiB)	Disks	Storage (GB)
fast-goose.maas		Ready		2	4	1	53.0
tangible-drum.maas		Deployed	ubuntu	4	8	2	107.0
thrifty-pipe.maas		Deployed	ubuntu	4	8	2	107.0
wasteful-riddle.maas		Deployed	ubuntu	8	15	1	107.0
wry-quilt.maas		Deployed	ubuntu	4	8	2	107.0

Data model

Machines	Networks	Users
Servers	L2	SSH
Devices	L3	Powerman
Switches*	Zones	API keys

Machine States

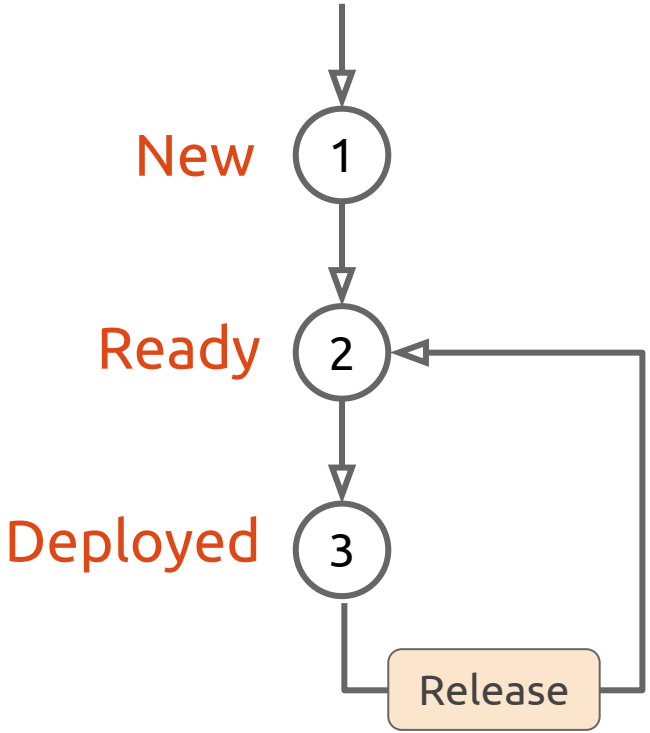

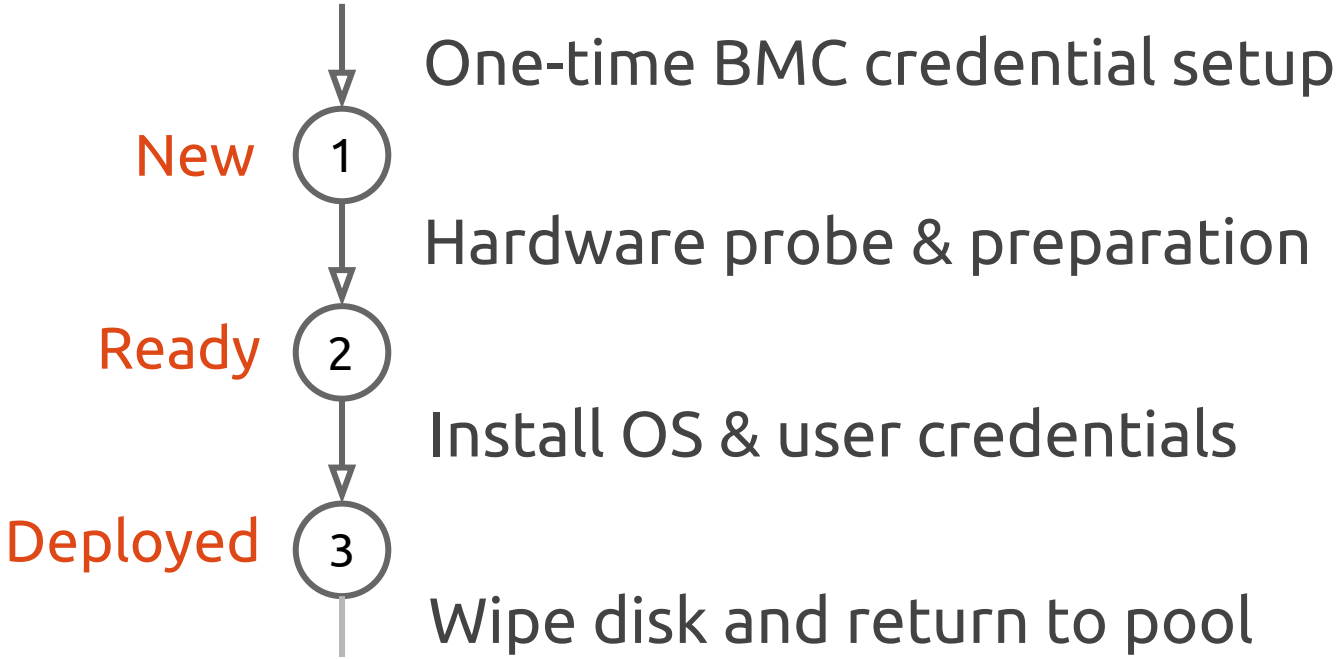
One-time BMC credential setup
New
1
Hardware probe & preparation
Ready
2
Install OS & user credentials
Deployed
3
Wipe disk and return to pool

IPMI Automatic Enlistment

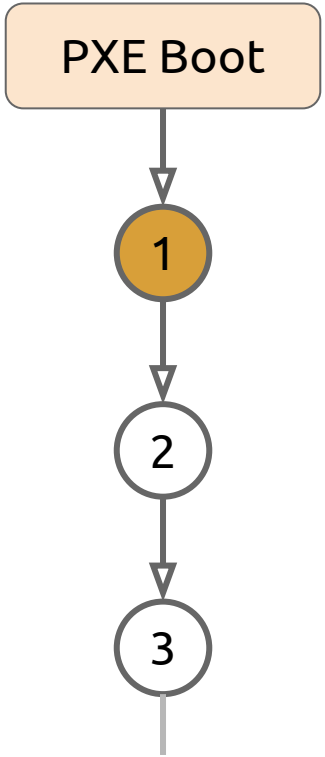

A. Boot into ephemeral environment

B. Set up credentials for IPMI in-band

C. Create random machine ID

D. Ready for remote control

Deployments are fast

HDD

8 minutes

SSD

8 minutes

Deployments are fast

Validated in the field

MAAS 2.0 Highlights

High availability

Scale-out architecture

Flexible network ranges

Custom DHCP configuration

Full DNS management

Faster disk wipe

OCP Hardware Tested w/ Ubuntu

Freedom, Windmill, Winterfell

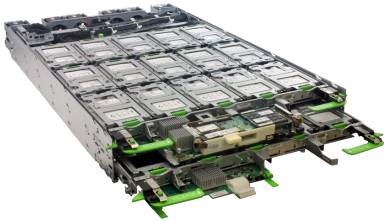

Roadrunner, Decathlete, Leopard,
Yosemite, Honey Badger, Panther, Knox

Microsoft Open Cloud Server (OCS)

Edgecore Networks AS5712-54X, AS 7712-32X,
Faceboook Wedge 40, Wedge 100

Open Compute
MAAS & Juju Deployments

UTSA
The University of Texas at San Antonio™

ITRI
Industrial Technology
Research Institute

清華大學
Tsinghua University

Labs

See it live

MAAS

Juju

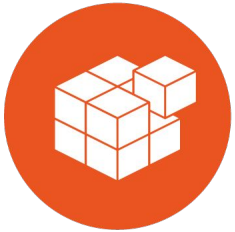

Ubuntu
Core

Christian Reis
kiko@canonical.com

David Duffey
david.duffey@canonical.com

Luke Williams
luke.williams@canonical.com

maas.io

Backup

MAAS Feature Highlights

User Interfaces

- Web (native)
- REST API
- CLI
- Python library

Operating Systems

- CentOS 6, 7
- RHEL 6,7
- Ubuntu
- Windows 2003, 2010
- Windows Nano*
- Custom*

Disk layout options

- Partitioned
- Raw
- RAID
- LVM
- bcache

NIC config options

- Raw
- Tagged VLAN
- Bonded
- Bridged(soon)

BMC protocols

- IPMI
- AMT
- Redfish(soon)

Chassis support

- Cisco UCS
- Microsoft OCS
- Seamicro 15K
- HP Moonshot

VM Controllers

- virsh for KVM & PowerKVM
- vSphere

Other

- SSH & Powerman credential setup
- Auto-enlistment via initial PXE-boot
- Node hardware inventory & auto-tagging
- Customization of commissioning and installation
- Dynamic allocation based on machine parameters

Today We ...

Analyzed Big Data with
Apache Hadoop, Zeppelin, Spark

Monitored a Kubernetes Cluster

Scaled and Installed OpenStack
(twice)

Deployed Snappy Ubuntu Core via
MAAS and ran FBOSS and SnapRoute
as Apps alongside other apps

Today You Can ...

Start Modeling @ jujucharms.com

Conjure-up openstack on your machine

Use Auto-pilot w/ MAAS

Try Snappy Ubuntu Core in the cloud

Installing Hadoop and Kubernetes

```
$ juju bootstrap


$ juju add-model bigdata                                    # optional!
$ juju deploy apache-hadoop-spark-zeppelin
$ juju list-actions spark
$ juju run-action spark/0 < sparkpi | pagerank | ... >


$ juju add-model kubernetes                                 # optional!
$ juju deploy observable-kubernetes
$ juju list-actions kibana
$ juju run-action kibana/0 load-dashboard dashboard=beats
```

Installing OpenStack at Facebook Labs

1. Install Ubuntu LTS
   $ sudo apt install maas juju
   $ sudo maas-region-admin createadmin # to set credentials

2. Setup MAAS @ http://<maas.ip>/MAAS/
   - Import images (click a button)
   - Tell MAAS which network(s) to manage (e.g. eth0/eth1)

3. Rack or Power on bare metal, and commission in MAAS

4. Point Juju to MAAS and deploy Big Software!
   $ juju bootstrap
   $ juju deploy openstack

Openstack Autopilot

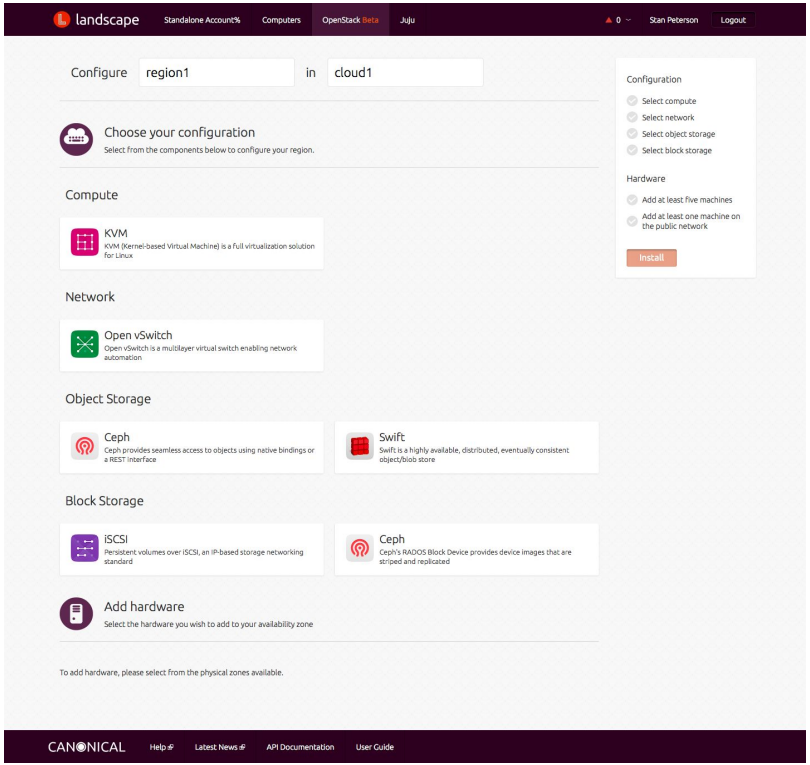

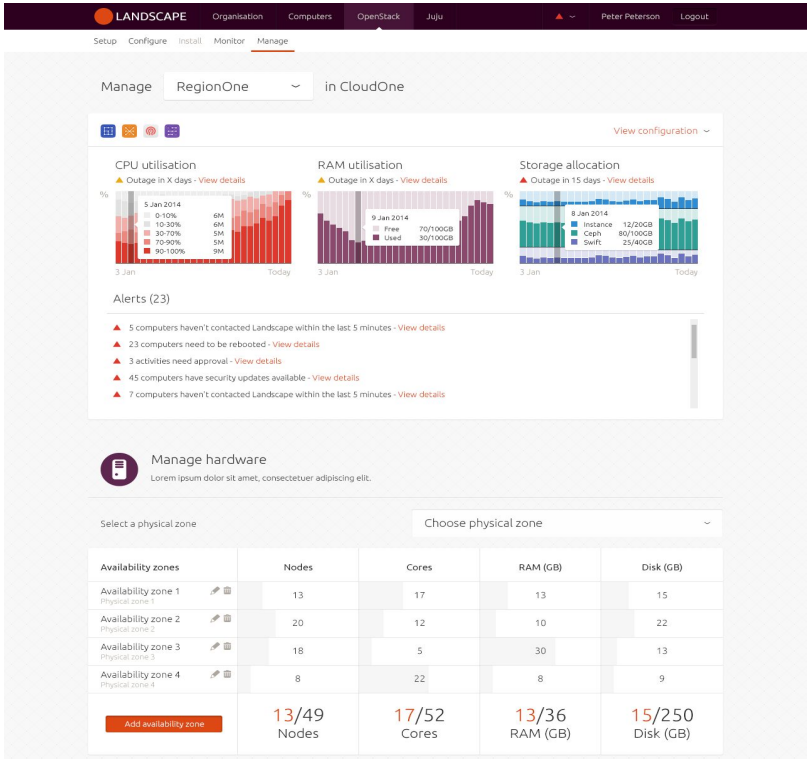

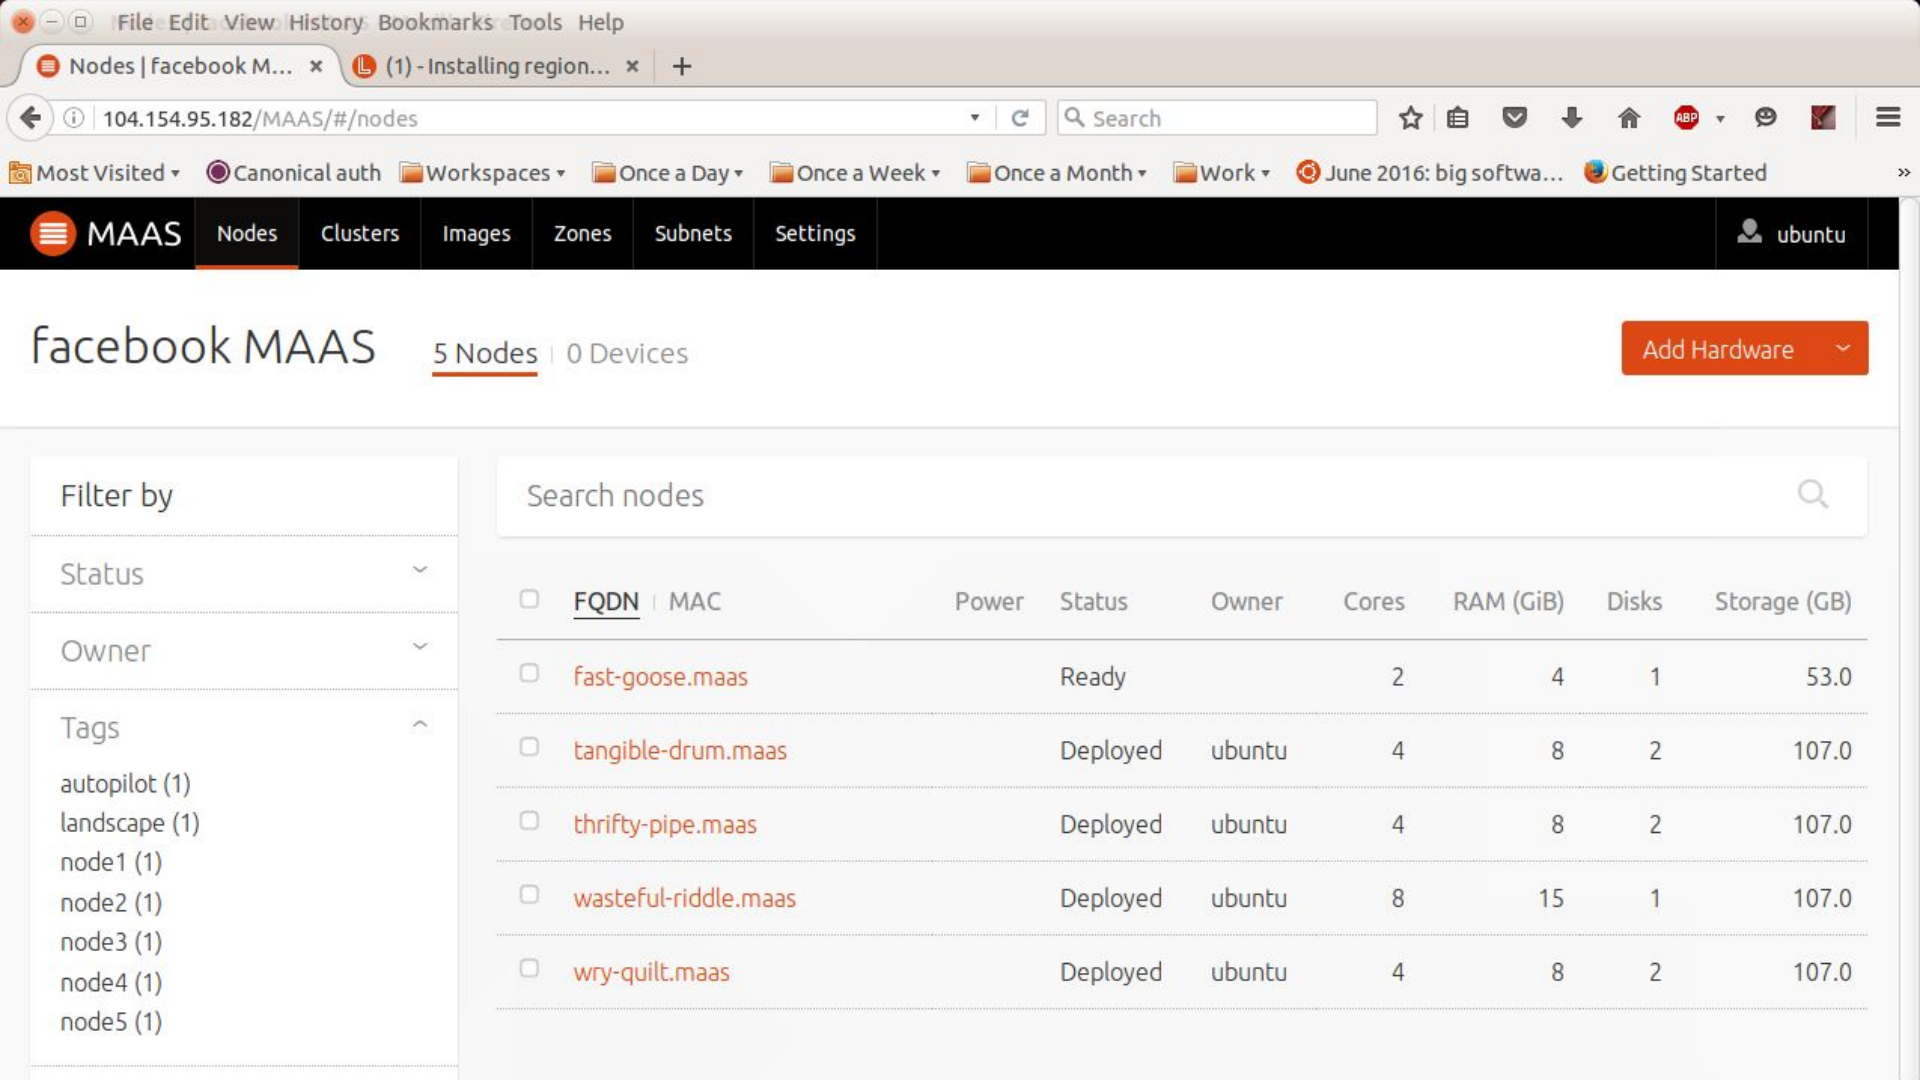
File Edit View History Bookmarks Tools Help
Nodes | facebook M...
(1) - Installing region...
104.154.95.182/MAAS/#/nodes
Search
Most Visited
Canonical auth
Workspaces
Once a Day
Once a Week
Once a Month
Work
June 2016: big softwa...
Getting Started
MAAS
Nodes
Clusters
Images
Zones
Subnets
Settings
ubuntu
facebook MAAS
5 Nodes
0 Devices
Add Hardware
Filter by
Status
Owner
Tags
autopilot (1)
landscape (1)
node1 (1)
node2 (1)
node3 (1)
node4 (1)
node5 (1)
Search nodes
FQDN
MAC
Power
Status
Owner
Cores
RAM (GiB)
Disks
Storage (GB)
fast-goose.maas
Ready
2
4
1
53.0
tangible-drum.maas
Deployed
ubuntu
4
8
2
107.0
thrifty-pipe.maas
Deployed
ubuntu
4
8
2
107.0
wasteful-riddle.maas
Deployed
ubuntu
8
15
1
107.0
wry-quilt.maas
Deployed
ubuntu
4
8
2
107.0

File Edit View History Bookmarks Tools Help

Nodes | facebook M... × (1) - Installing region... ×

https://10.20.100.111/account/standalone/openstack/cloud/2/region/2/install

Search

Most Visited ▾ Canonical auth Workspaces ▾ Once a Day ▾ Once a Week ▾ Once a Month ▾ Work ▾ June 2016: big softwa... Getting Started

landscape | Standalone Account | Computers | OpenStack | Juju | 1 | ubuntu | Logout

Installing menlo in facebook

93%

Status	Description	Created at
In progress	Wait for relations to establish	Today 14:13 UTC
Waiting	Wait for relations to finish establishing	Today 14:13 UTC
Waiting	Create a Landscape admin user	Today 14:13 UTC
Waiting	Configure availability zones	Today 14:13 UTC
Waiting	Create the public network	Today 14:13 UTC
Waiting	Create network for user	Today 14:13 UTC
Waiting	Create an OpenStack account for david.duffey@canonical.com	Today 14:13 UTC
Waiting	Create network for user	Today 14:13 UTC

menlo in facebook installed

100%

Deployment is complete. The region is now ready to use.

Manage your region

Get the most from your cloud

Take advantage of Canonical's cloud tools and services.

Ubuntu Advantage

Juju

Spread the word

File Edit View History Bookmarks Tools Help

Nodes | facebook M... | (1) - Manage menlo i...

https://10.20.100.111/account/standalone/openstack/cloud/2/region/2/

Search

Most Visited | Canonical auth | Workspaces | Once a Day | Once a Week | Once a Month | Work | June 2016: big softwa... | Getting Started

landscape | Standalone Account | Computers | OpenStack | Juju | 1 | ubuntu | Logout

Manage menlo in facebook

CPU utilisation

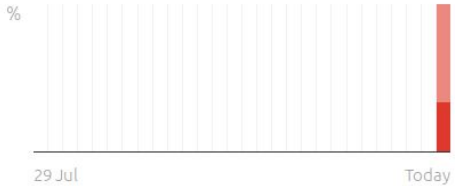

RAM utilisation

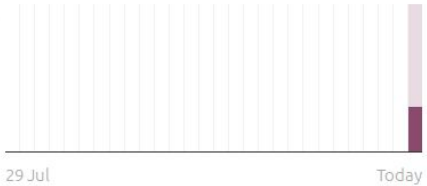

Storage allocation

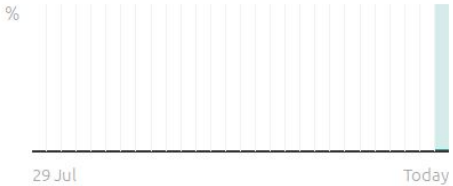

Region configuration

Configuration

- KVM
- Open vSwitch
- Ceph
- Ceph

Hardware

- 1 availability zone(s)
- 3 nodes
- 12 cores
- 23.4GiB RAM
- 150GB Disk

Next actions

Download OpenStack RC file for david.duffey@canonical.com

Download Juju configuration for david.duffey@canonical.com

Download OpenStack RC file for admin

Remove | OpenStack Dashboard

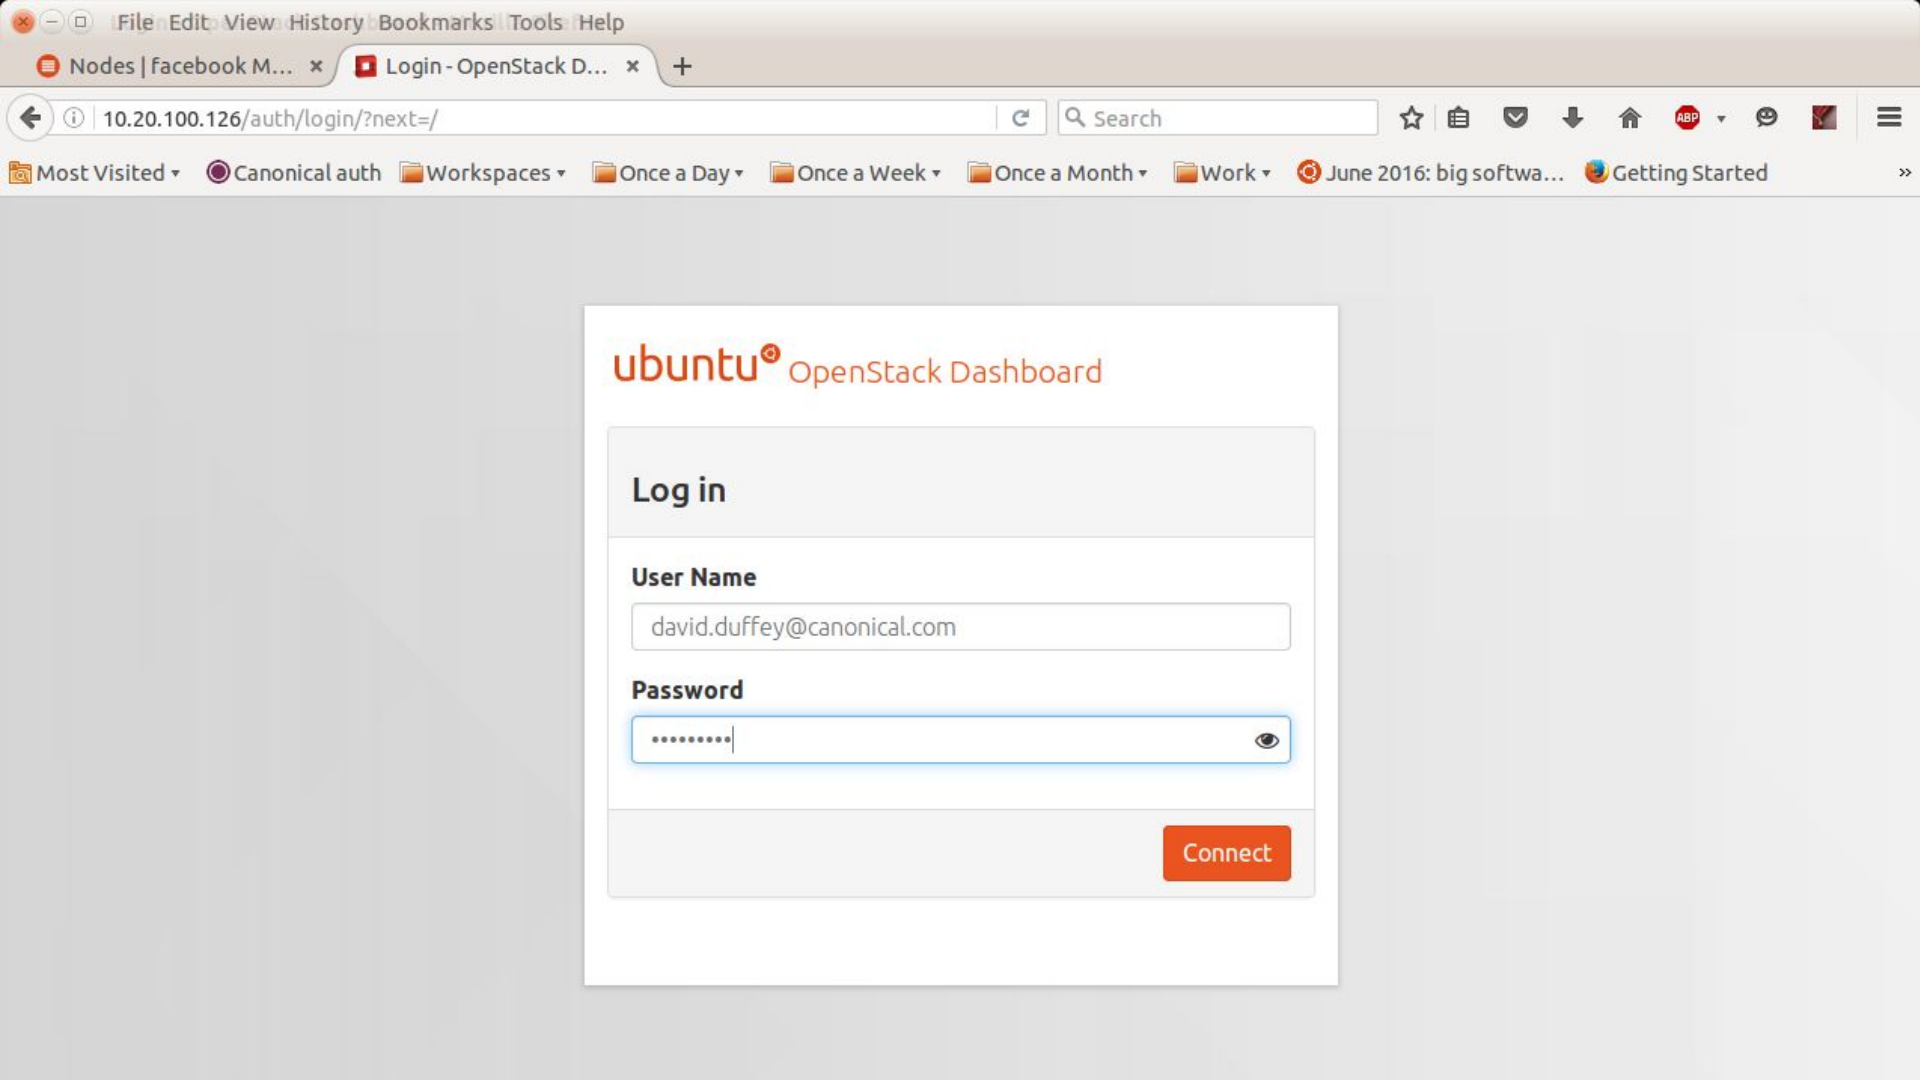
File Edit View History Bookmarks Tools Help
Nodes | facebook M...
Login - OpenStack D...
10.20.100.126/auth/login/?next=/
Search
Most Visited
Canonical auth
Workspaces
Once a Day
Once a Week
Once a Month
Work
June 2016: big softwa...
Getting Started
ubuntu OpenStack Dashboard
Log in
User Name
david.duffey@canonical.com
Password
Connect

File Edit View History Bookmarks Tools Help

Nodes | facebook M... | Instance Overview - ...

10.20.100.126/project/

Search

Most Visited | Canonical auth | Workspaces | Once a Day | Once a Week | Once a Month | Work | June 2016: big softwa... | Getting Started

ubuntu | david_duffey_project | david.duffey@canonical.com

- Project
- Compute
  - Overview
  - Instances
  - Volumes
  - Images
  - Access & Security
- Network
- Object Store
- Identity

Overview

Limit Summary

Instances
Used 0 of 100

VCPUs
Used 0 of 200

RAM
Used 0 of 512,000

Floating IPs
Used 0 of 100

Security Groups
Used 1 of 100

Volumes
Used 0 of 10

Volume Storage
Used 0 of 1,000

Usage Summary

File Edit View History Bookmarks Tools Help

Nodes | facebook M... | Instances - OpenSta...

10.20.100.126/project/instances/

Search

Most Visited | Canonical auth | Workspaces | Once a Day | Once a Week | Once a Month | Work | June 2016: big softwa... | Getting Started

ubuntu | david_duffey_project | david.duffey@canonical.com

- Project
- Compute
  - Overview
  - Instances
  - Volumes
  - Images
  - Access & Security
- Network
- Object Store
- Identity

Instances

Instance Name = | Filter | Launch Instance | Delete Instances | More Actions

Instance Name	Image Name	IP Address	Size	Key Pair	Status	Availability Zone	Task	Power State	Time since created	Actions
test	auto-sync/ubuntu-xenial-16.04-amd64-server-20160815-disk1.img	10.10.0.4	m1.small	-	Active	nova	None	Running	4 minutes	Create Snapshot

Displaying 1 item

Snaps and Ubuntu Core

Transactionality, security and app stores

Snaps are Universal Linux Packages

Security, Transactionality, App Stores

Package any app for every Linux desktop, server, cloud or device, and deliver updates directly.

Learn to craft snaps or join the community

Read the installation instructions or build snapd from source

Disaggregating the Network Operating System

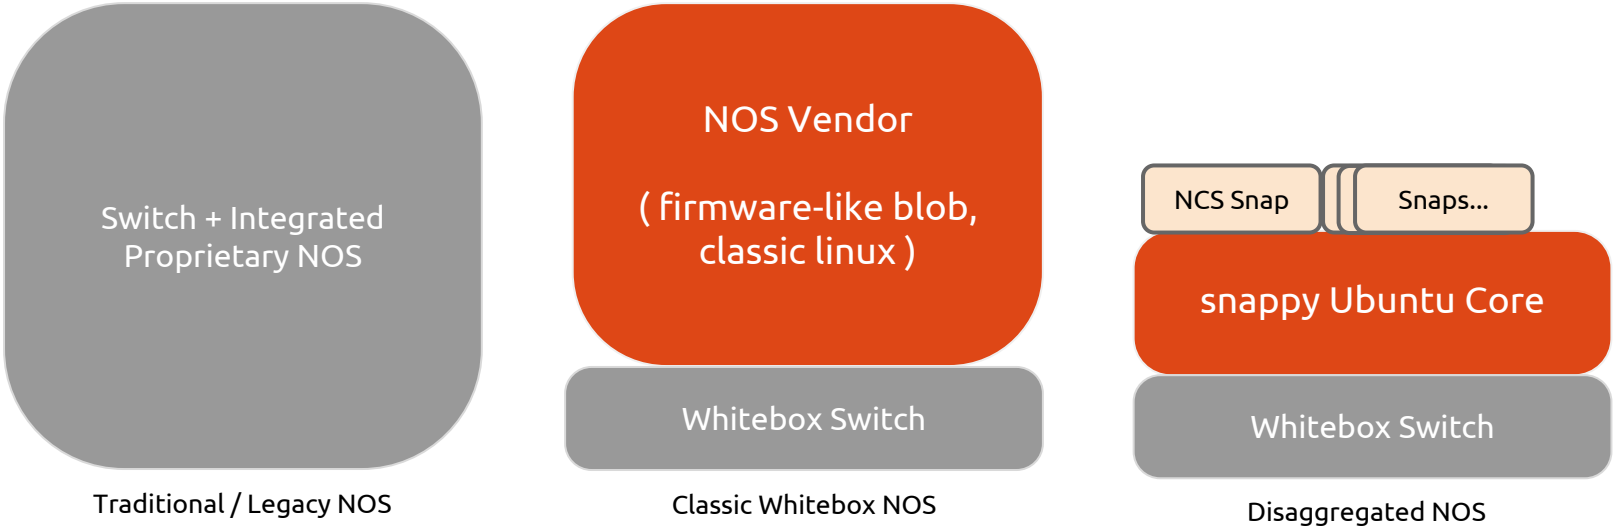

Traditional / Legacy NOS

Classic Whitebox NOS

Disaggregated NOS

classic Linux

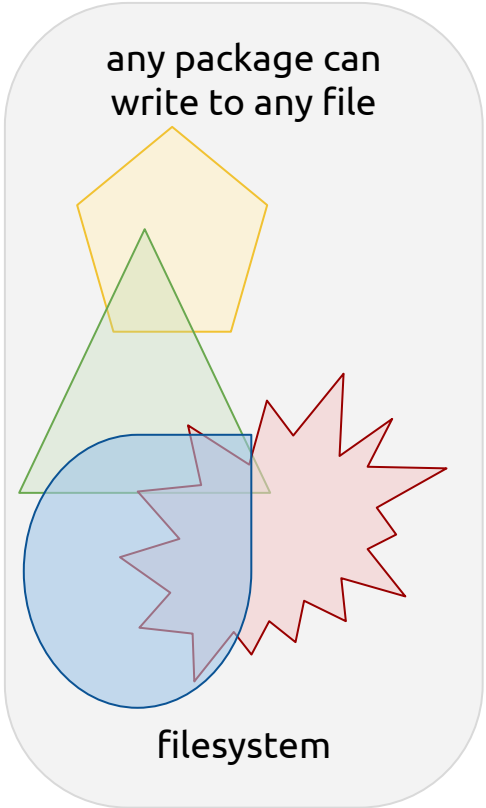

snappy Ubuntu Core

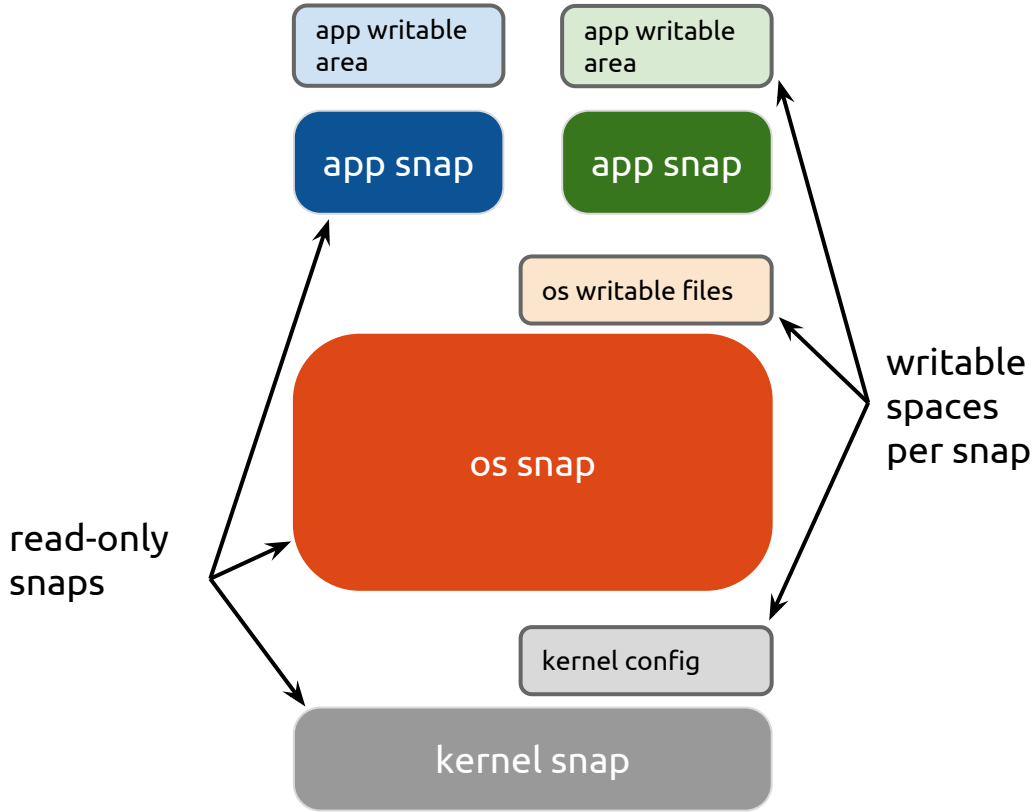

Transactional update and rollback

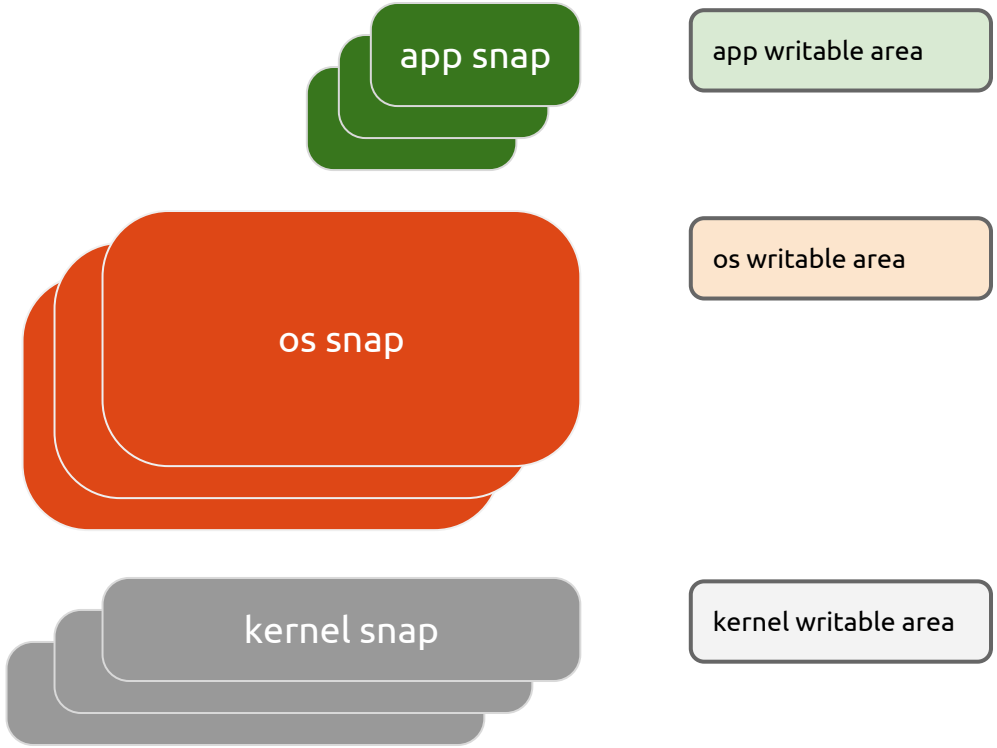

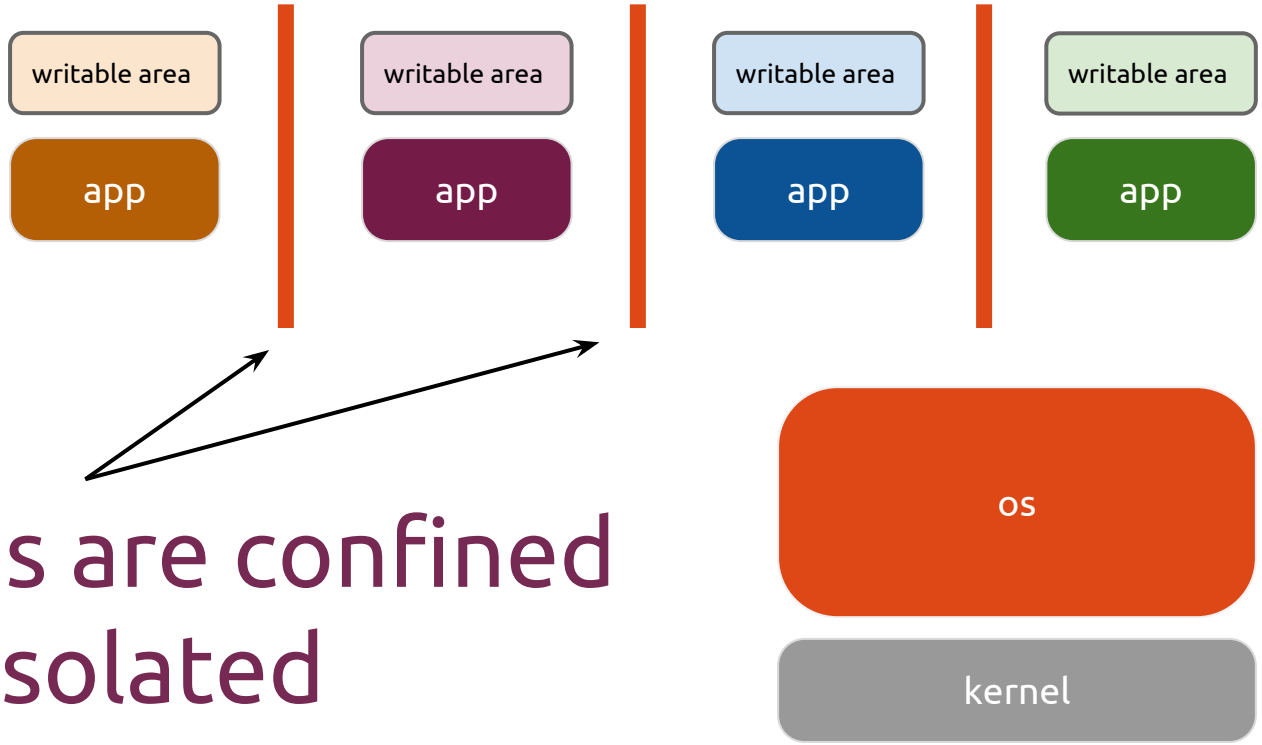

Snaps are confined and isolated

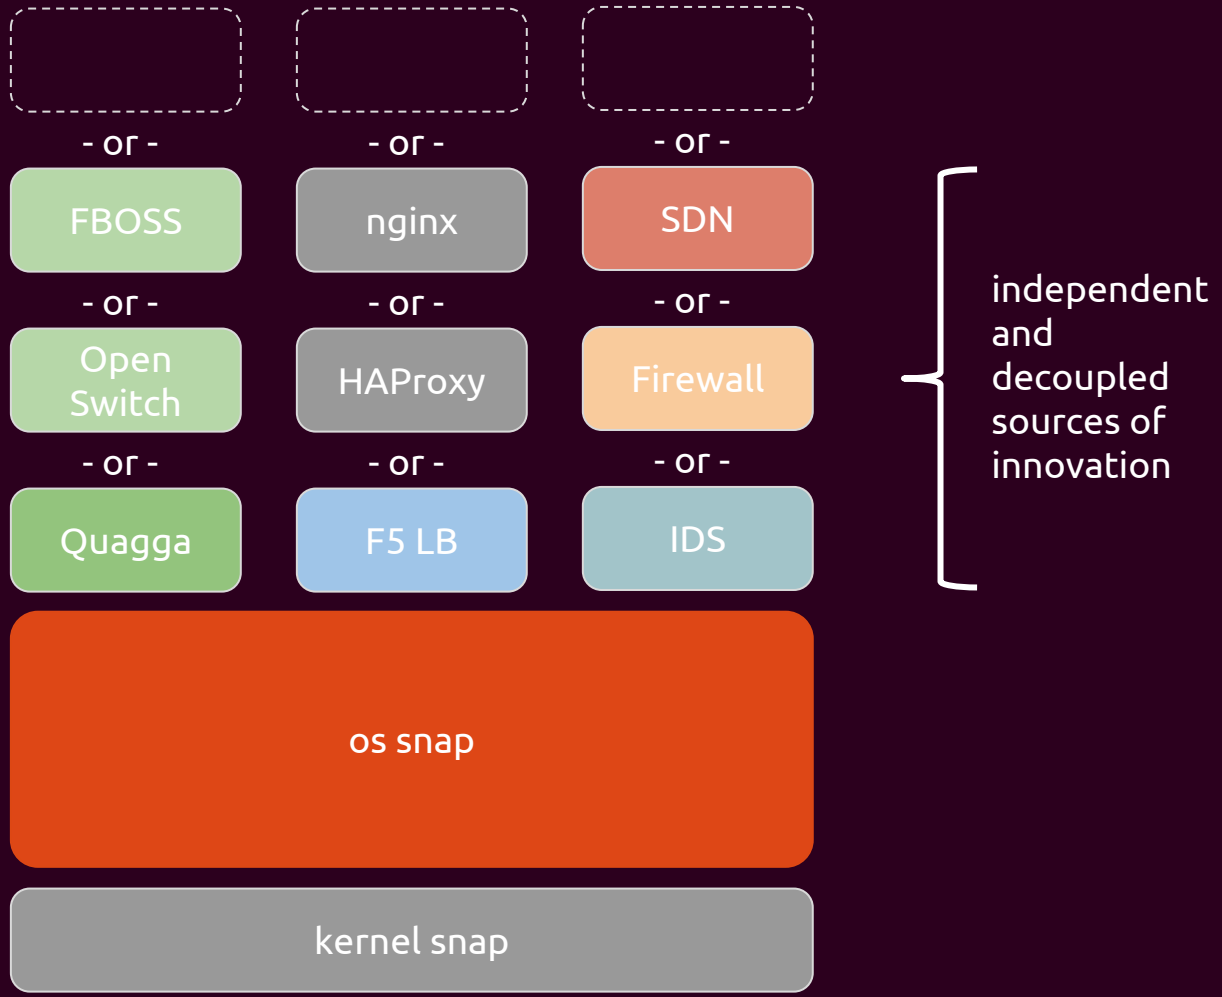
- or -
- or -
- or -
FBOSS
nginx
SDN
- or -
- or -
- or -
Open Switch
HAProxy
Firewall
- or -
- or -
- or -
Quagga
F5 LB
IDS
os snap
kernel snap
independent and decoupled sources of innovation

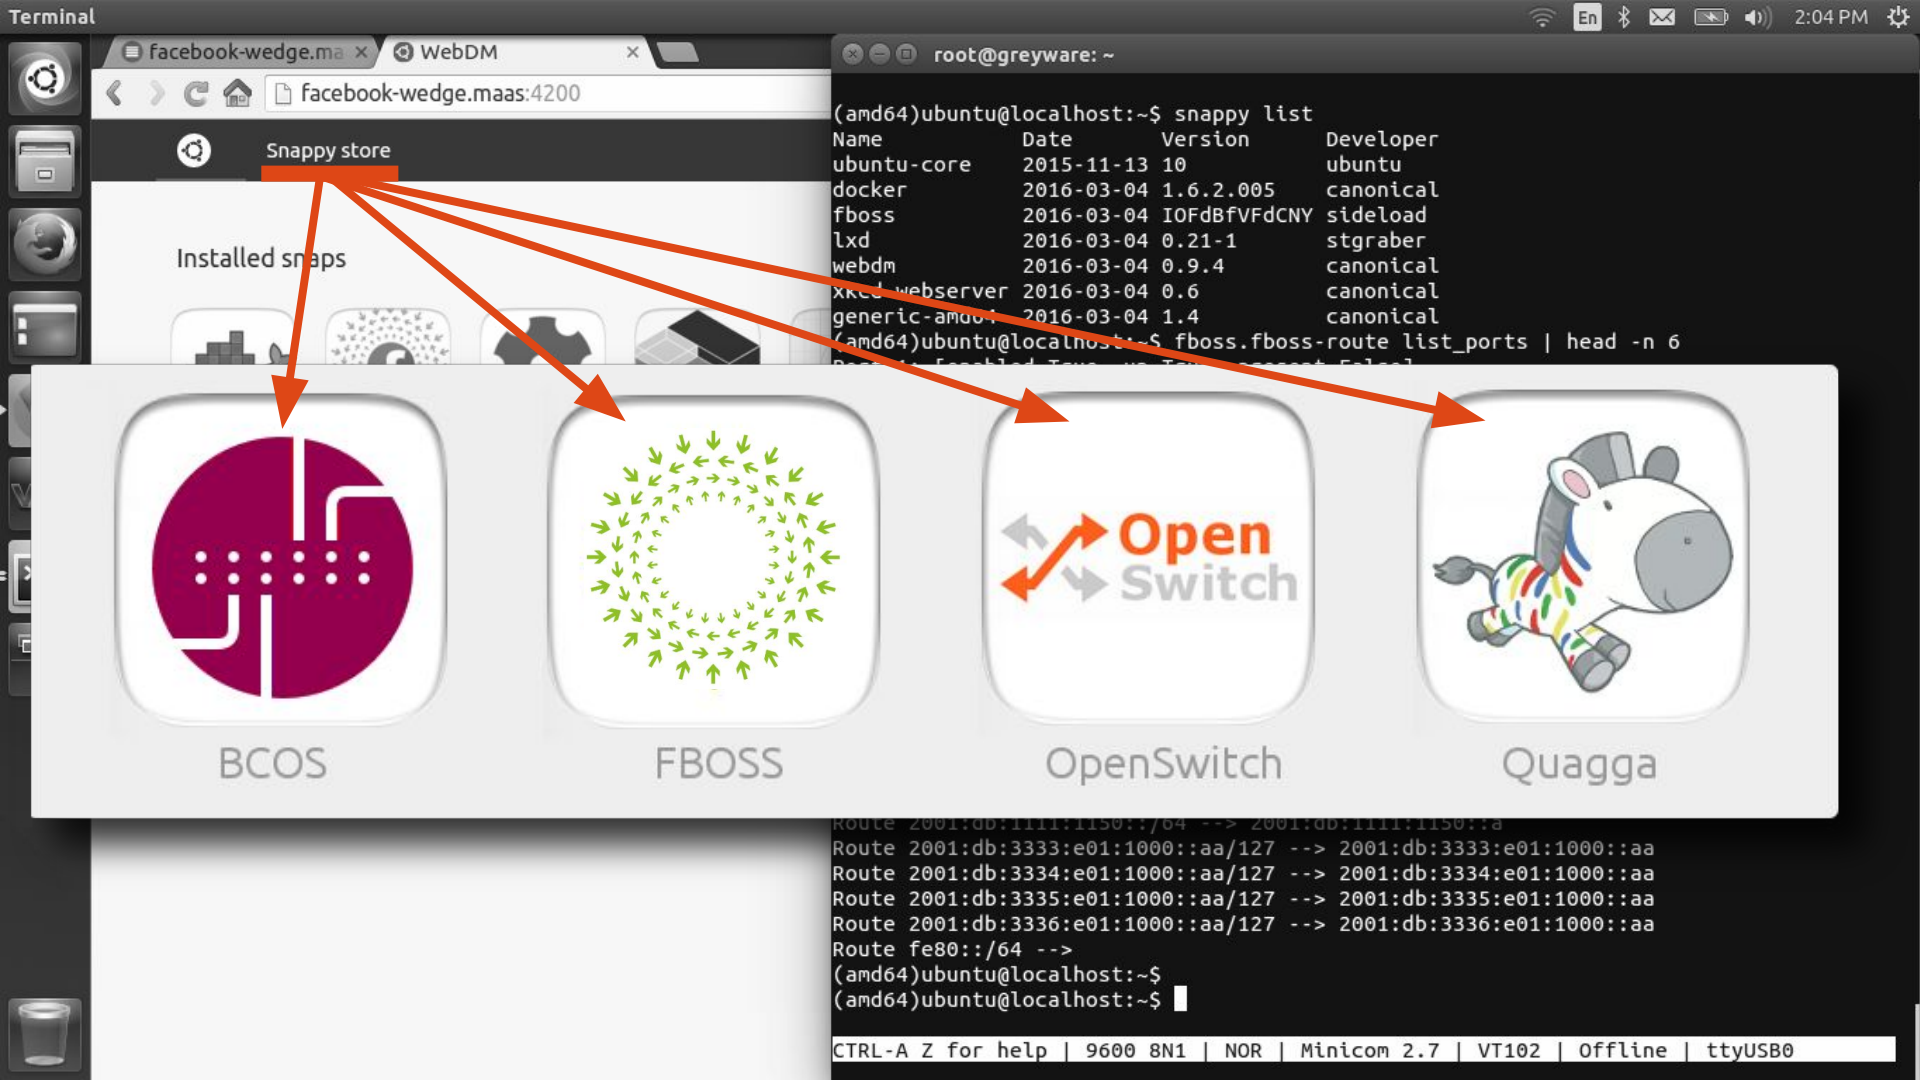
Terminal
2:04 PM
facebook-wedge.ma
WebDM
facebook-wedge.maas:4200
Snappy store
Installed snaps
root@greyware: ~
(amd64)ubuntu@localhost:~$ snappy list
Name Date Version Developer
ubuntu-core 2015-11-13 10 ubuntu
docker 2016-03-04 1.6.2.005 canonical
fboss 2016-03-04 IOFdBfVFdCNY sideload
lxd 2016-03-04 0.21-1 stgraber
webdm 2016-03-04 0.9.4 canonical
webserver 2016-03-04 0.6 canonical
generic-amd64 2016-03-04 1.4 canonical
(amd64)ubuntu@localhost:~$ fboss.fboss-route list_ports | head -n 6
Route 2001:db:3333:e01:1000::aa/127 --> 2001:db:3333:e01:1000::aa
Route 2001:db:3334:e01:1000::aa/127 --> 2001:db:3334:e01:1000::aa
Route 2001:db:3335:e01:1000::aa/127 --> 2001:db:3335:e01:1000::aa
Route 2001:db:3336:e01:1000::aa/127 --> 2001:db:3336:e01:1000::aa
Route fe80::/64 -->
(amd64)ubuntu@localhost:~$
(amd64)ubuntu@localhost:~$
CTRL-A Z for help | 9600 8N1 | NOR | Minicom 2.7 | VT102 | Offline | ttyUSB0
Open
Switch
BCOS
FBOSS
OpenSwitch
Quagga

Demo of Data Centre Staging and Operations

MAAS

HA Region Controller

Bare-Metal-As-A-Service, running as a Snap

- Hardware inventory
- OS provisioning
- PXE
- IPMI
- DHCP
- DNS

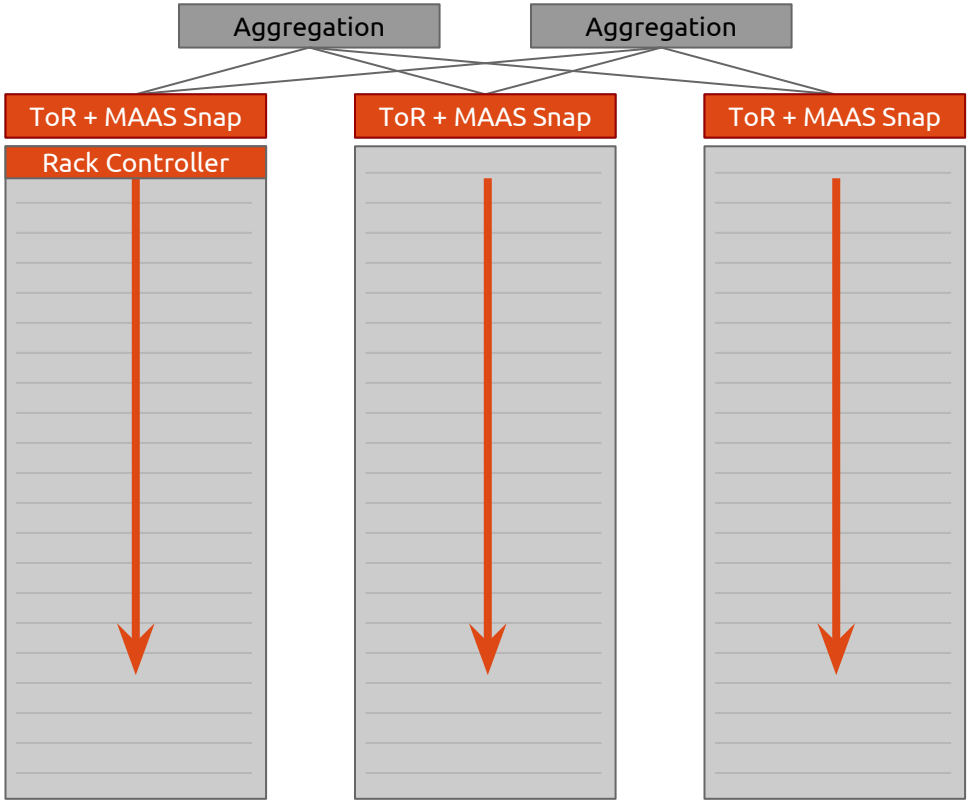

snappy

snappy extensible, run more Snap Apps

- Disaster Recovery
- Data Analytics
- Network Tools

The Perfect Switch OS

snappy
Ubuntu Core

Snappy Market
for extensibility

Ubuntu Core with
snappy **transactional updates**

Applications are confined by Canonical's
AppArmor kernel security system

reliable, worry free **updates** with rollback

Amazing application install experience
with **snappy packages "Snaps"**

Easily extensible with the **snappy store**

Perfect for **container** workloads

Ubuntu Core for Cloud Development

Microsoft Azure

Snappy images are in public clouds like AWS, Azure, and GCE

Vagrant and KVM images are ideal for rapid prototyping

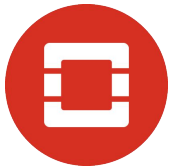

Snappy is available for private clouds

snappy

ubuntu.com/snappy