

From Bare Metal to Cloud

Andy Edmonds, @dizz, ICCLab, ZHAW
Piotr Kasprzak, GWDG

Intros

ICCLab

- Zurich University for Applied Sciences
- Cloud Computing Research



GWDG

- Service Provider for Max Planck Society and University of Goettingen
- Research



We've Hardware for Cloud!



GWDG Cloud Hardware


Nodes	38
CPUs	152
Core	2432
Memory	9728 TB

ICCLab Cloud Hardware

Nodes	20
CPUs	80
Core	1280
Memory	1920 TB

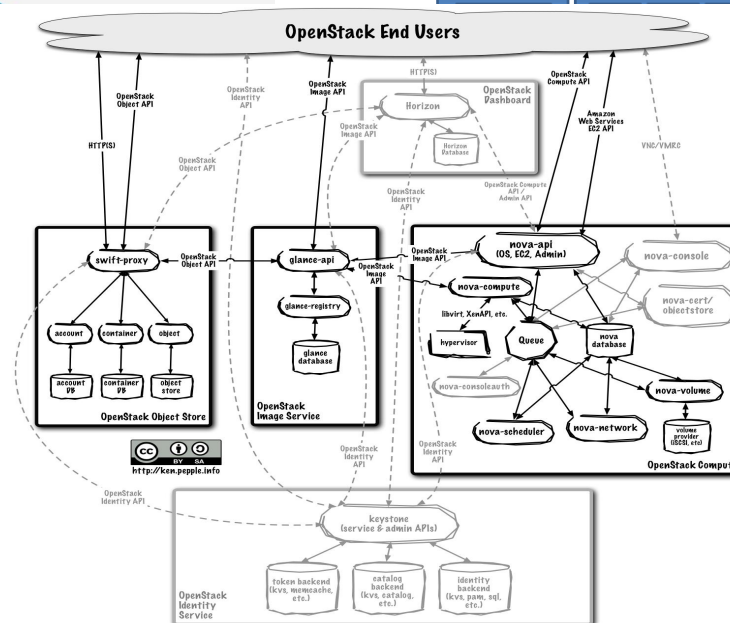
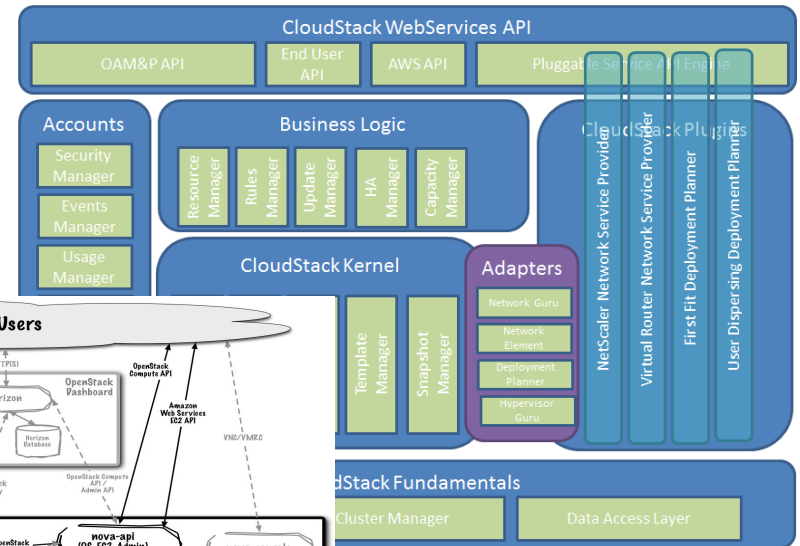
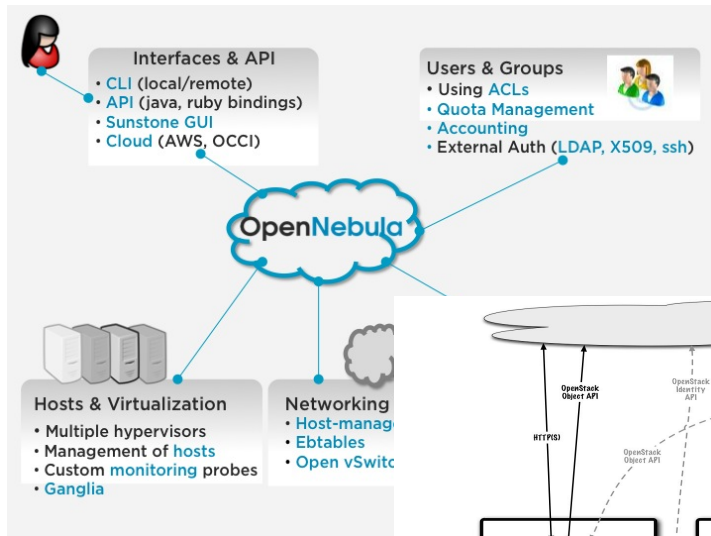
Challenges or Problems?

- Clouds in essence are big data centres
 - Means lots of servers:
 - Manual configuration **not an option**
 - Automation is **required**



Challenges or Problems?

Cloud frameworks **can/are be complicated!**



Challenges or Problems?

- But Clouds are "cool" - Aayyy!

BUT

- How to deploy a "cloud"
 - with minimal user interaction?
 - least number of "hands"?
 - across many servers?

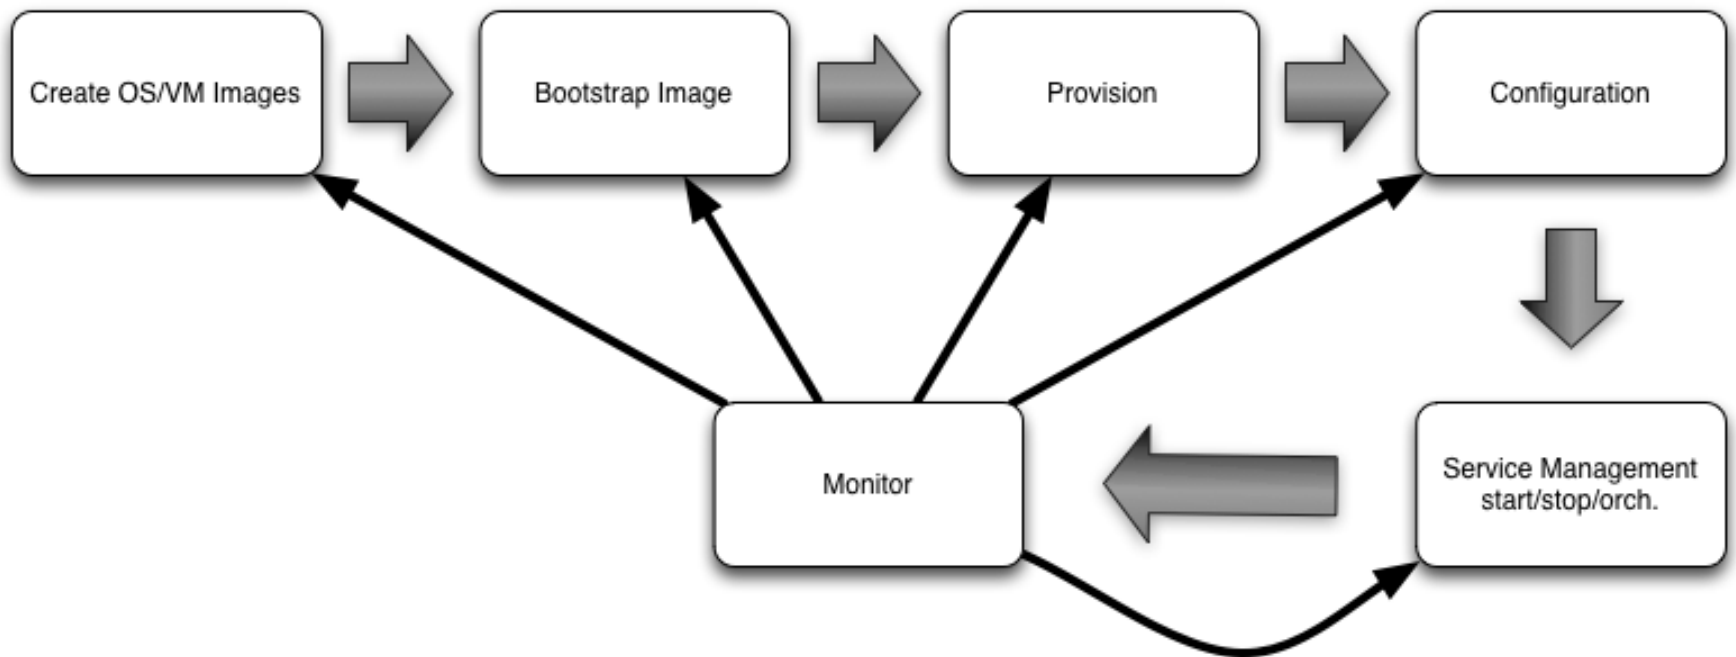


Challenges or Problems?

- How to **share/standardise** these processes?
 - Configuration - drift prevention
 - Testing - configuration, system functionality
 - Compliance - auditing, ITIL
 - Agility
 - Independence
 - Of physical/virtual deployment
 - Of infrastructure



Automation Toolchain



Automation Toolchain



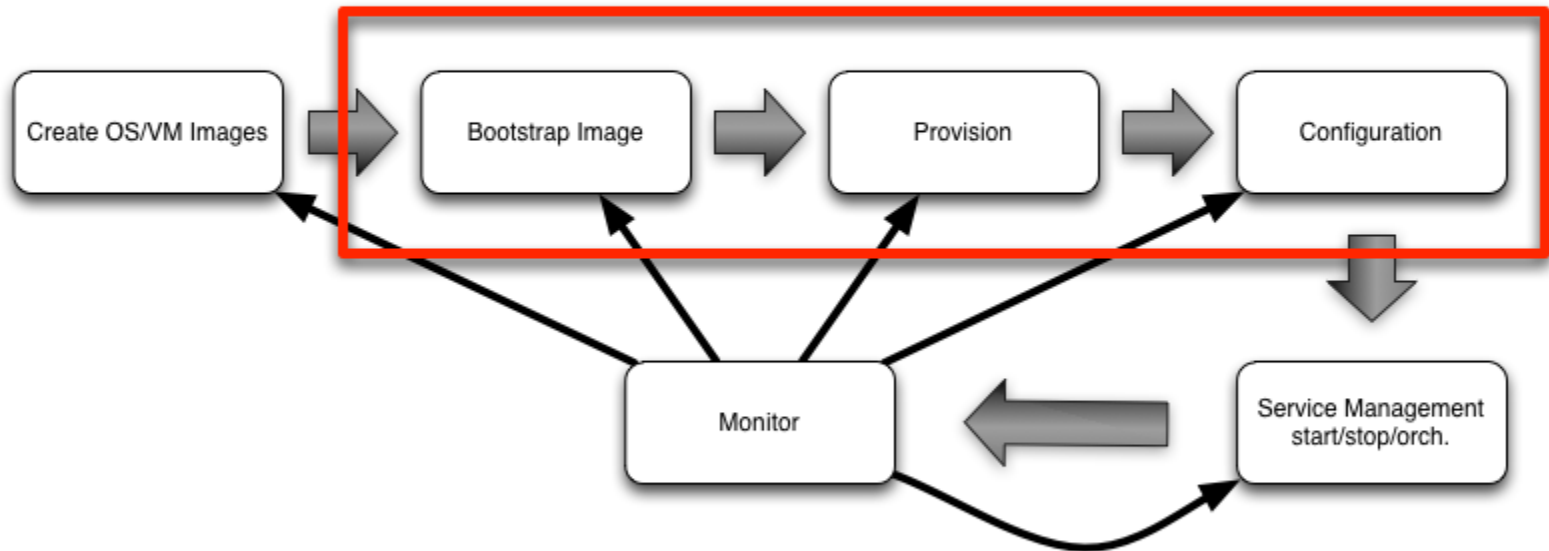
Ubuntu Netboot
apt-cache-ng




Foreman




Puppet



Provision - OS rollout



fedora^f



Baremetal

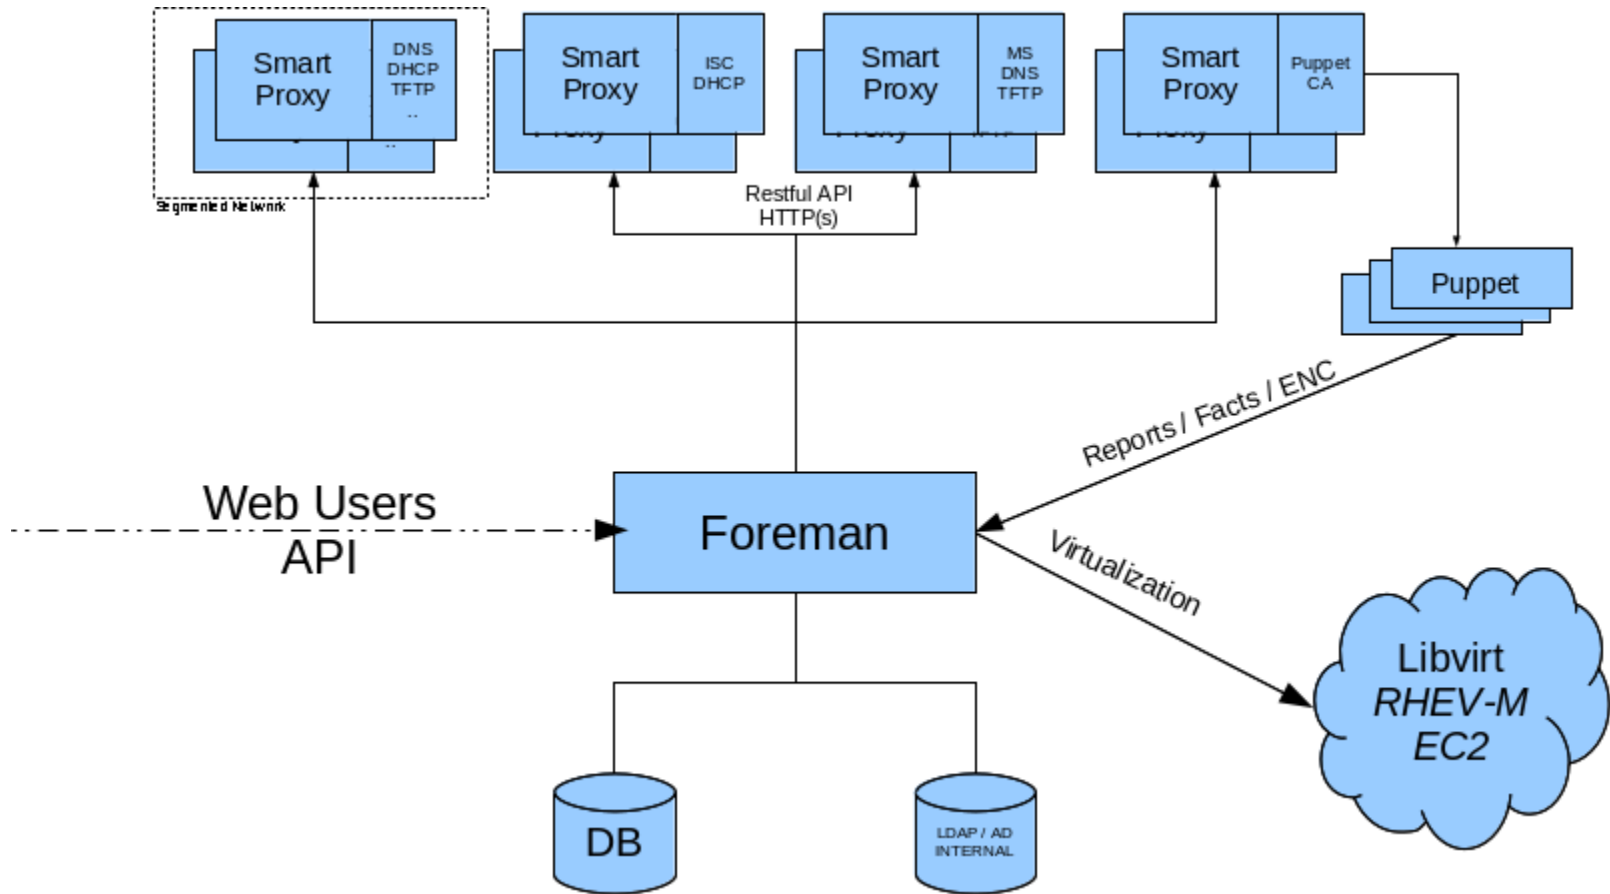
VM

Provision - Foreman



- "Single Address For All Machines Lifecycle Management".
- Manages or proxies to DNS, DHCP, TFTP, Virtual Machines, PuppetCA, CMDB
- Integrates with Puppet (and acts as web front end to it).
- Provisions:
 - most flavours of *NIX, Windows
 - Virtual machines - libvirt, oVirt
 - Cloud Resources - Amazon EC2, VMware vCenter
- Has an API! :-)

Provision - Foreman Arch

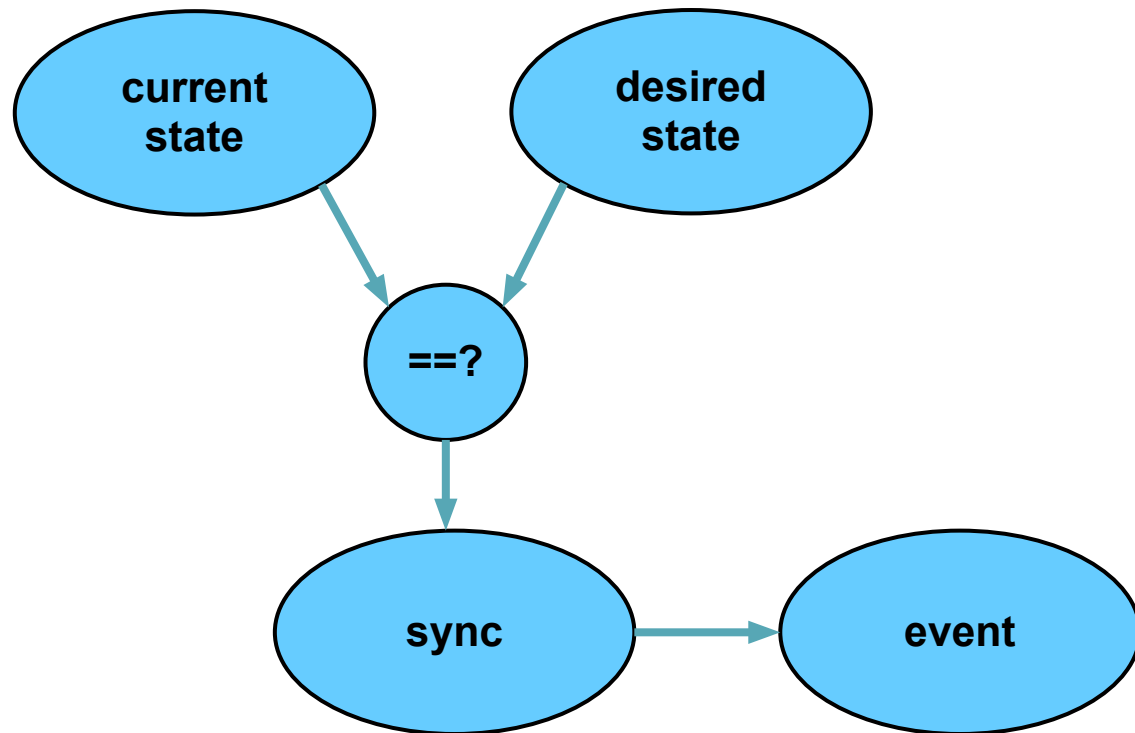


Configuration - Puppet

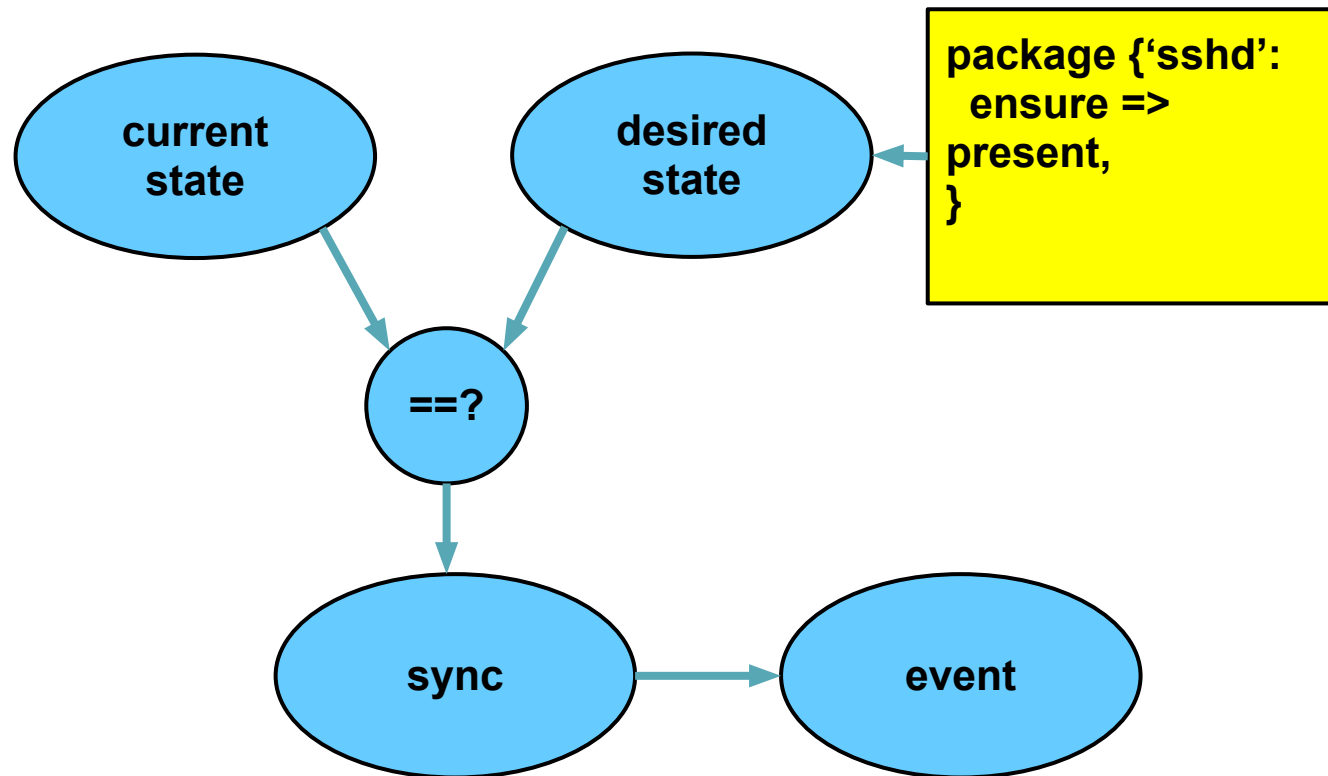


- Declarative configuration language
 - Describe desired state of a system, not how to achieve it
 - Idempotence
- Different types of resources: software package, service, user, configuration file, mysql database, ...
- Dependencies can be formulated
- Grouping of resources by "class" concept:
 - Way of structuring your descriptions
- Abstraction layer for resources:
 - Independence from system type (different variants of linux, *bsd, mac os, windows, ...)

Configuration - Puppet's Model



You describe system state...



Puppet collects current state...



```
rpm -q sshd
```

```
-----  
dpkg-query -  
search sshd
```



ubuntu

current
state

desired
state

```
package {'sshd':  
  ensure =>  
  present,  
}
```

==?

sync



event



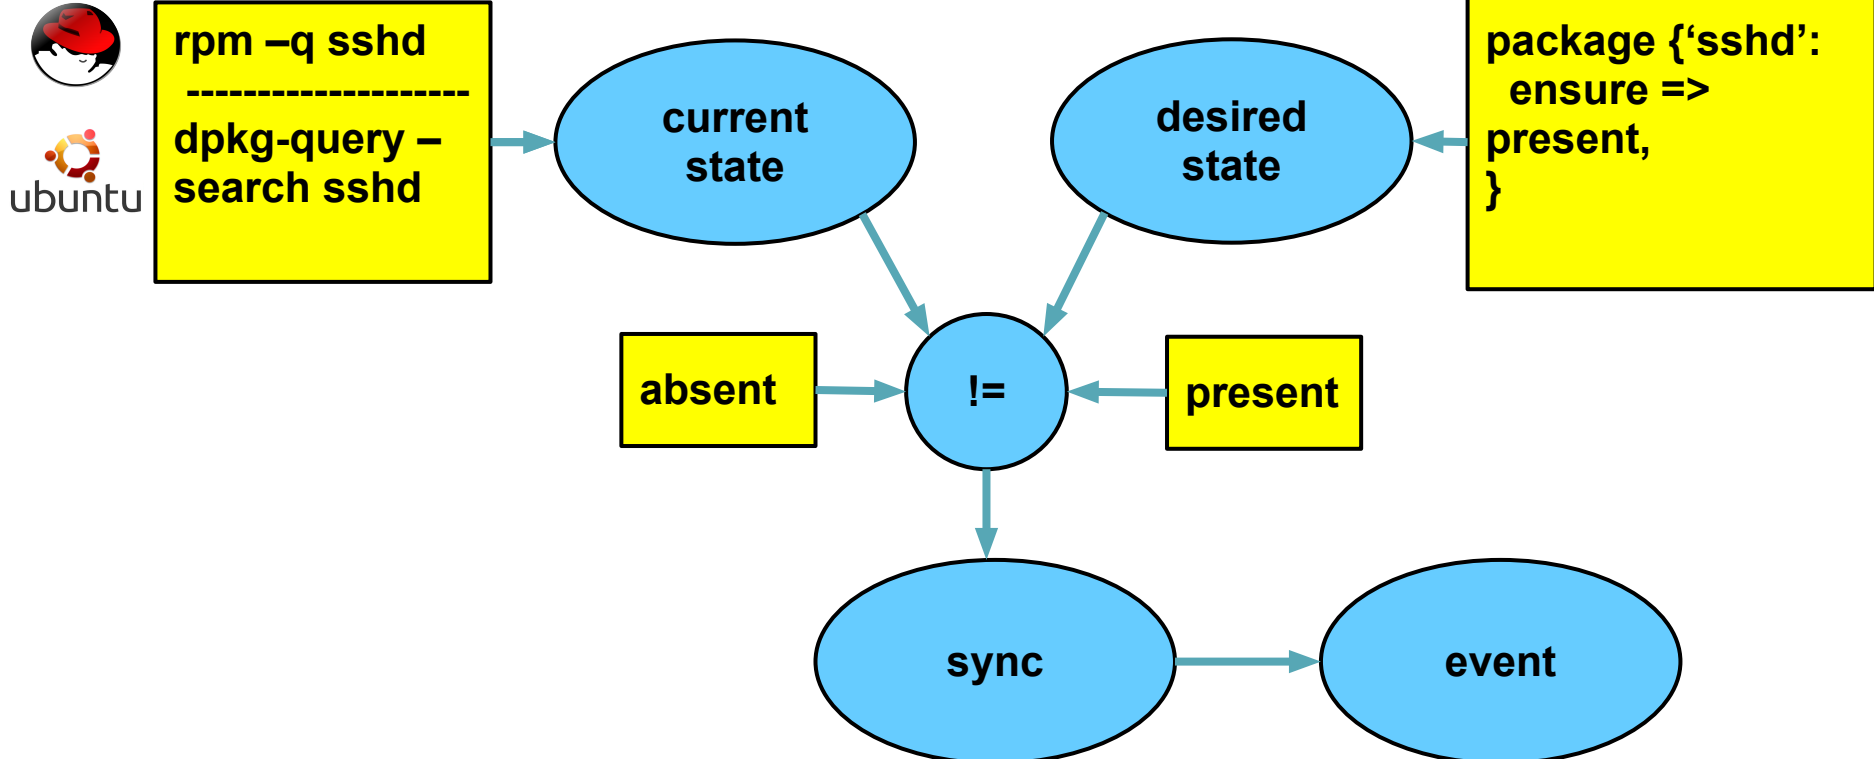



Puppet compares...



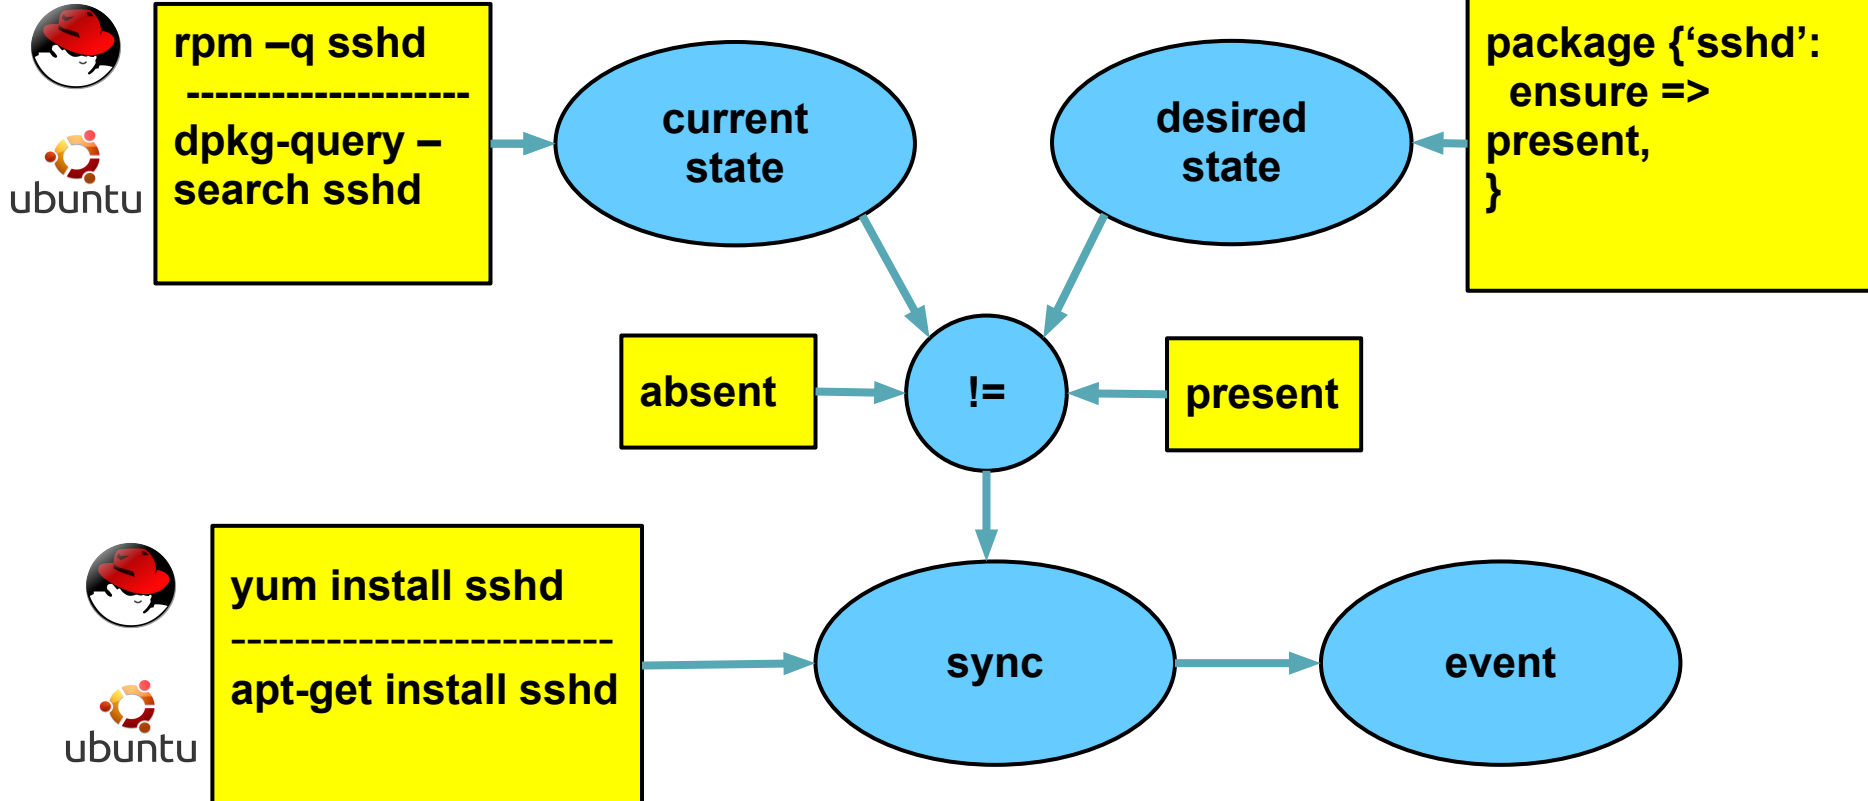
```
rpm -q sshd
-----
dpkg-query -
search sshd
```



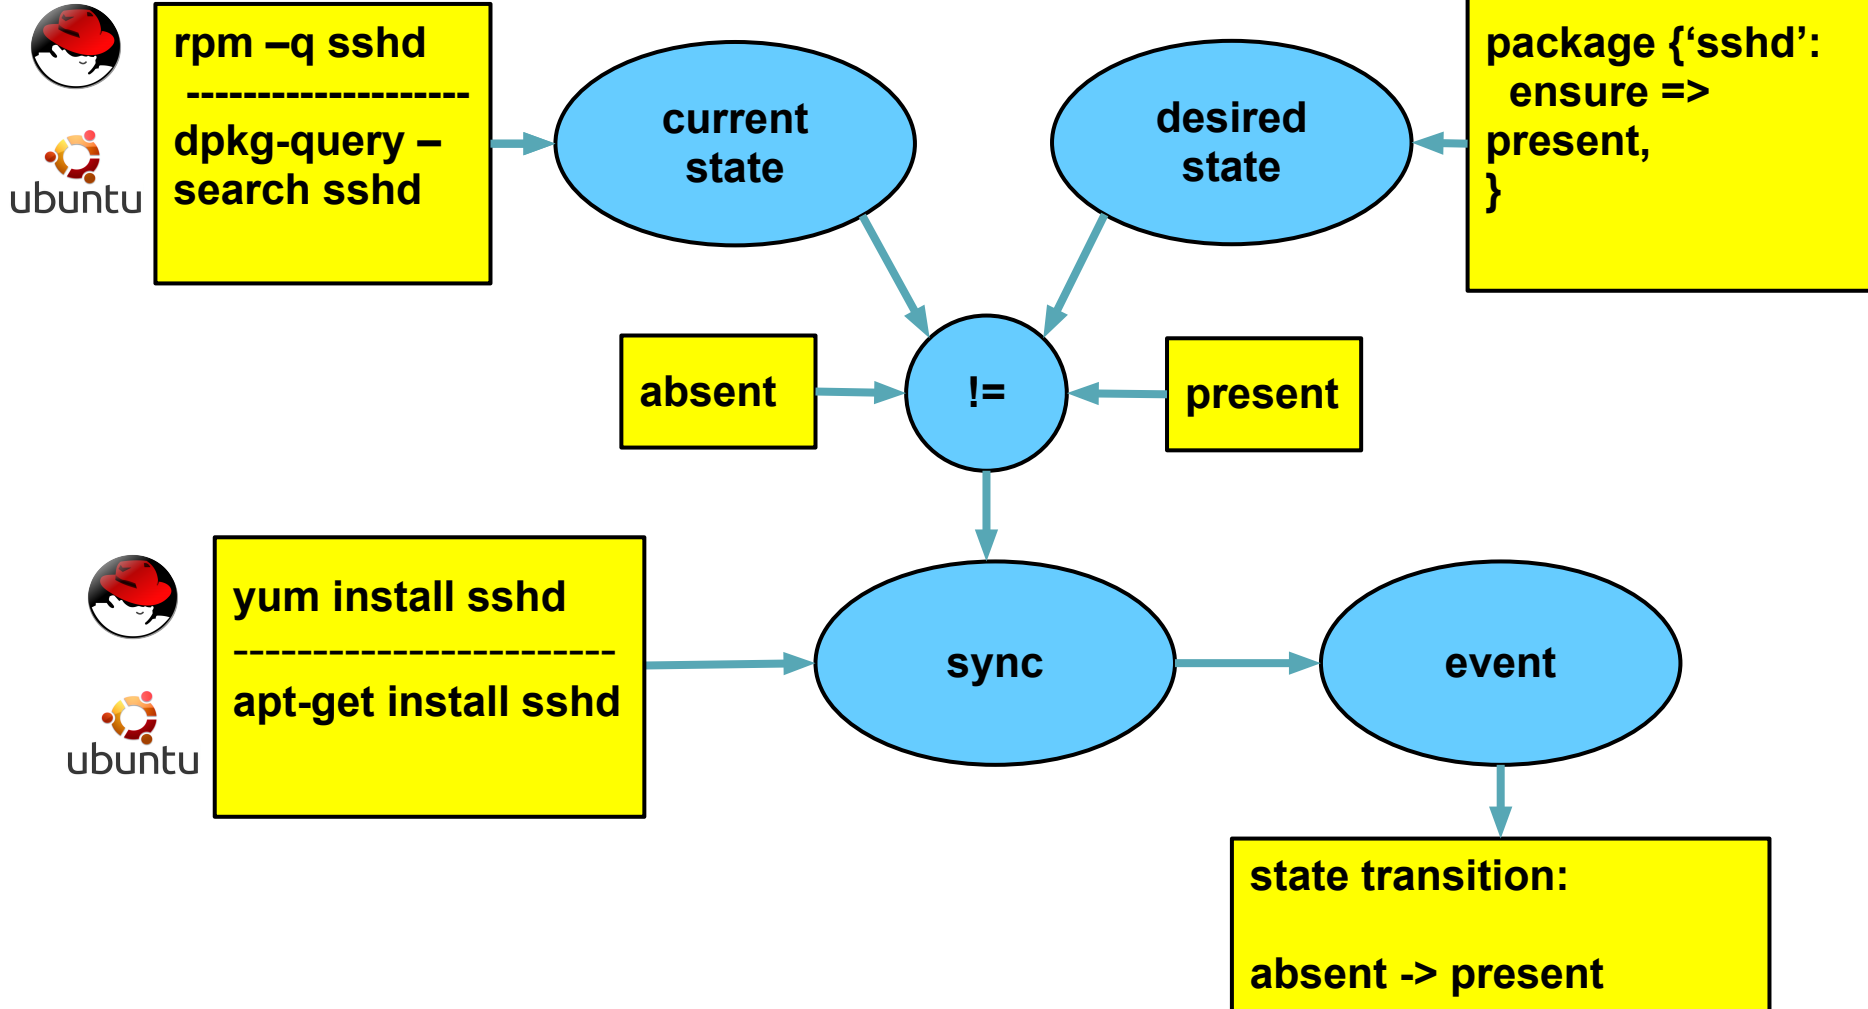
```
package {'sshd':
  ensure =>
  present,
}
```



Puppet synchronizes...



Puppet logs...



A more complete puppet manifest



```
class ssh::install {
  package { "openssh":
    ensure => present, }
}
class ssh::config {
  file { "/etc/ssh/sshd_config":
    ensure    => present,
    owner     => 'root',
    group     => 'root',
    mode      => 0600,
    source    => "puppet:///modules/ssh/sshd_config",
    require   => Class["ssh::install"],
    notify    => Class["ssh::service"], }
}
class ssh::service {
  service { "sshd":
    ensure      => running,
    hasstatus   => true,
    hasrestart  => true,
    enable      => true,
    require     => Class["ssh::config"], }
}
class ssh {
  include ssh::install, ssh::config, ssh::service
}
```



dependency
"if I change..."

OpenStack @ 10,000m, Looks Easy!

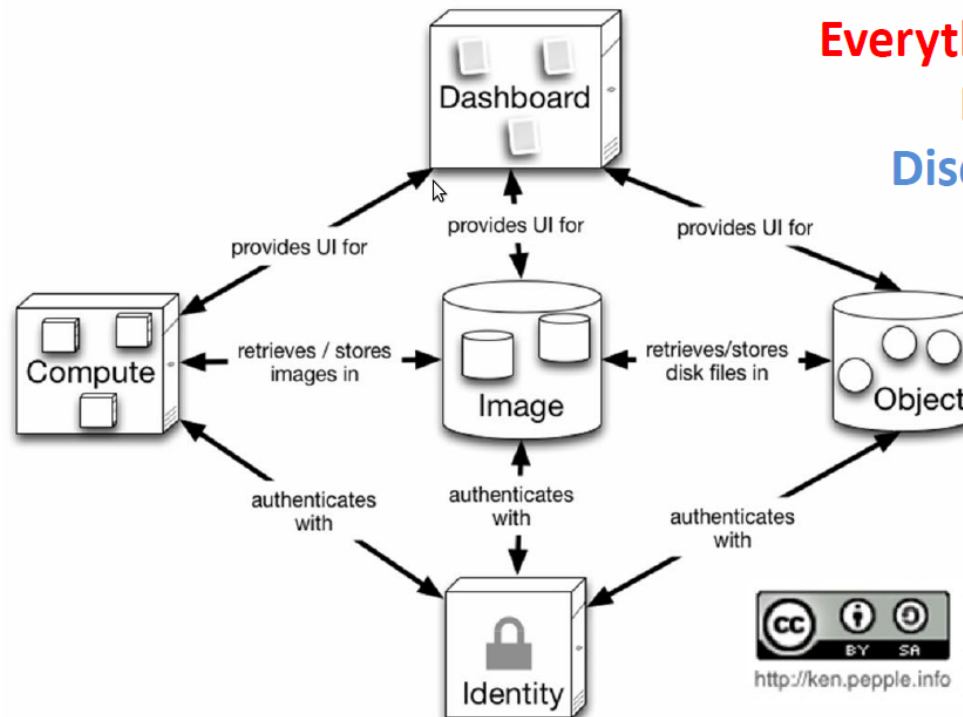
Software as a Service

Platform as a Service

Infrastructure as a Service



openstack™



Everything has an API
Message based
Discrete Pluggable
Components



<http://ken.pepple.info>

OpenStack - The Ugly Close-up

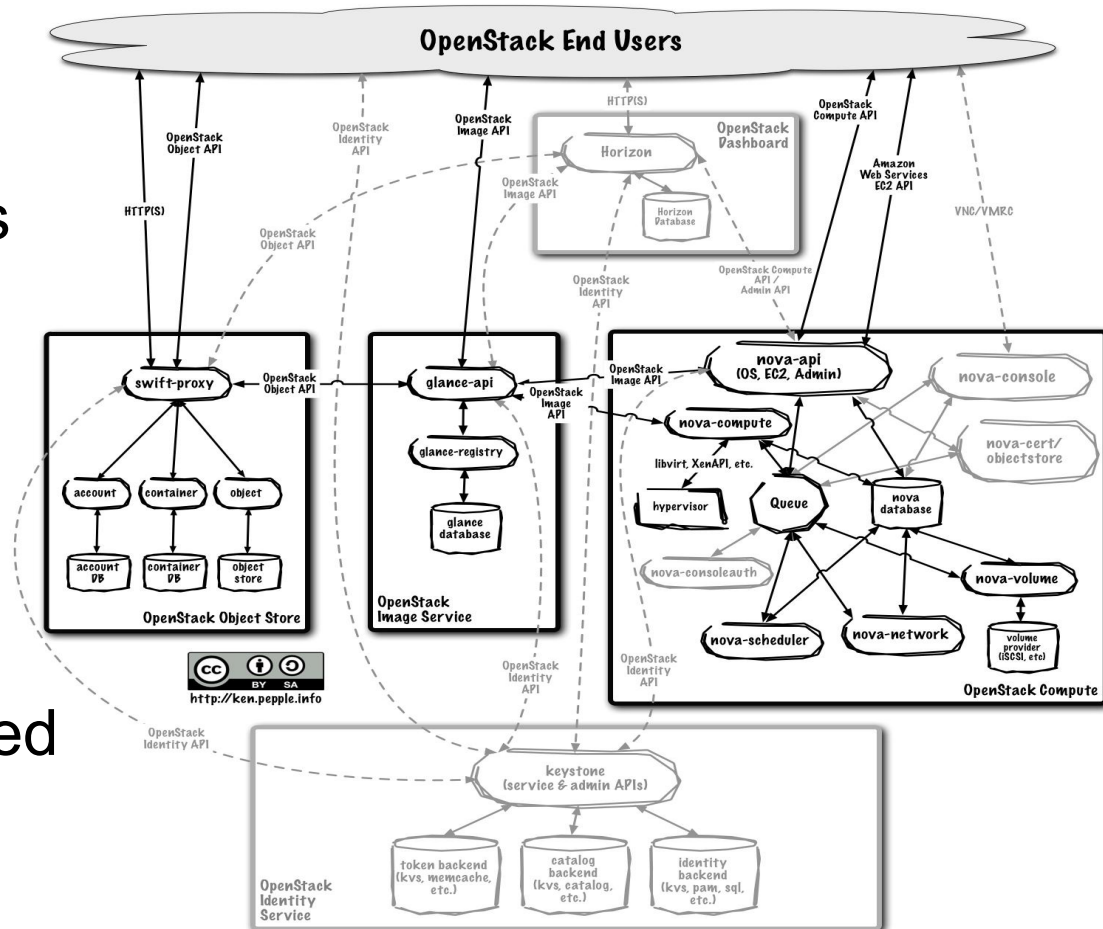
Complicated

- Many Services
- Many Dependencies

Challenge to deploy

- 100's, 1000's of nodes?

You **need** an automated toolchain!



Apple Moment!



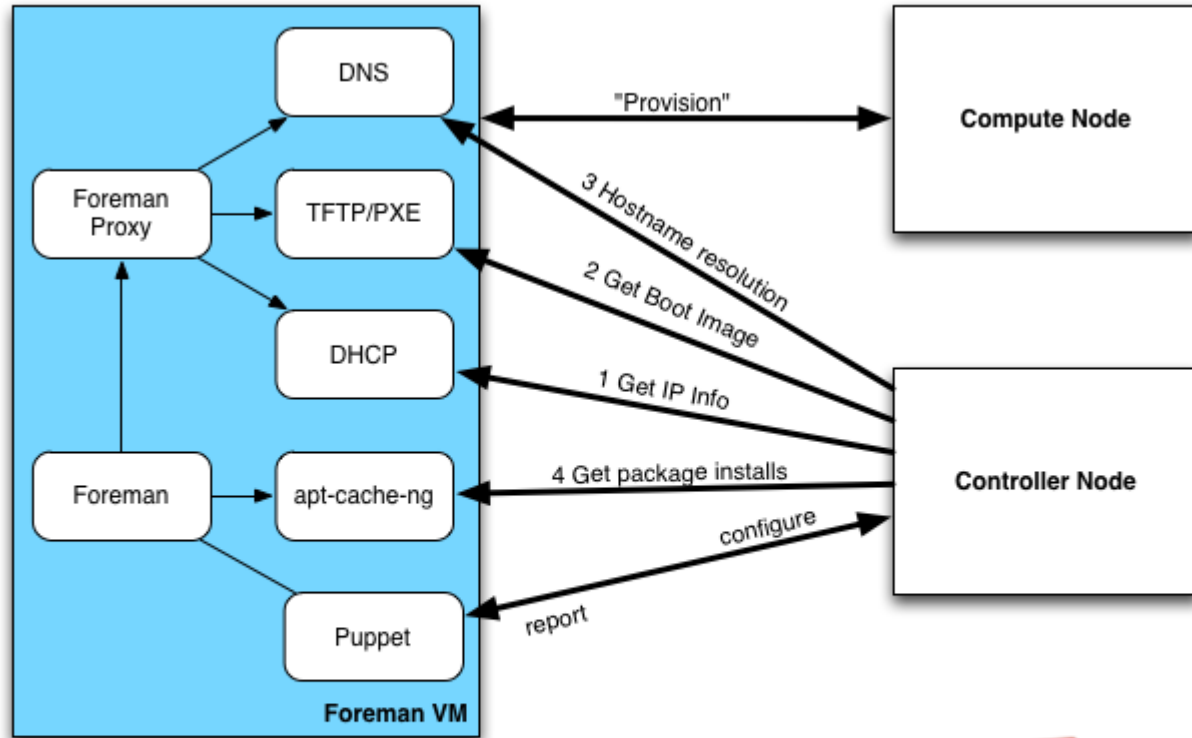
Demo - What could go wrong?!

Multi-node OpenStack Installation

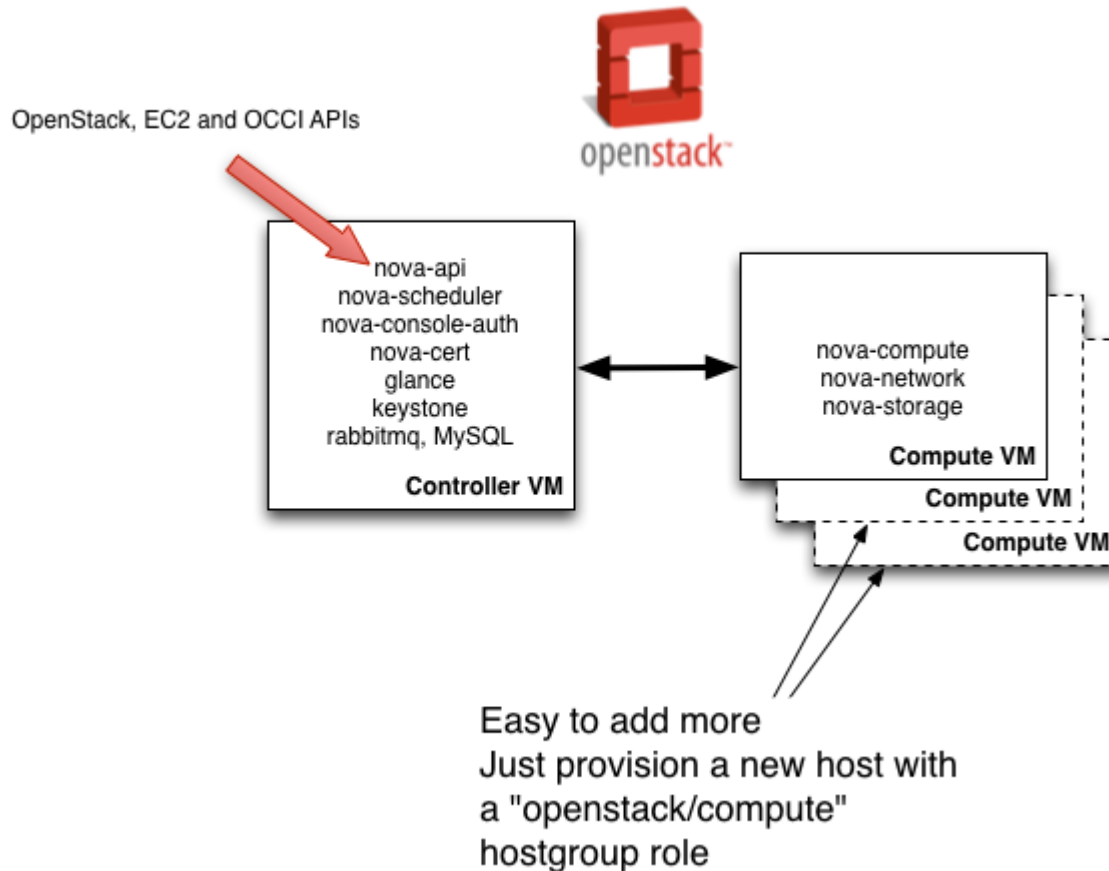
- 1 controller node
 - "boss"
- 1 compute node
 - "worker1"
- More time? Easy to add more.



Demo: Deployment Architecture



Demo: OpenStack Component Deployment



Demo: Code/Config Details

- There are 2 roles (*hostgroups*)
 - **openstack/controller** - `controller.pp`
 - **openstack/compute** - `compute.pp`
- Both have different puppet manifests
 - Same 'icclab' module

```
root@foreman:/etc/puppet/modules/iaas# tree icclab
icclab
├── manifests
│   ├── all_in_one.pp
│   ├── controller.pp
│   ├── compute.pp
│   └── params.pp
└── 1 directory, 4 files
```

What's in a controller node?

```
1 class icclab::controller{
2
3     include icclab::params
4
5     $admin_password      = 'admin_pass'
6     $keystone_admin_token = 'keystone_pass'
7
8     class { 'openstack::controller':
9
10        public_address      => $icclab::params::controller_node_public,
11        public_interface    => $icclab::params::public_interface,
12        private_interface   => $icclab::params::private_interface,
13        internal_address    => $icclab::params::controller_node_internal,
14        floating_range      => '192.168.56.128/25',
15        fixed_range         => $icclab::params::fixed_range,
16        multi_host          => true,
17        network_manager     => $icclab::params::network_manager,
18        verbose             => true,
19        auto_assign_floating_ip => false,
20        mysql_root_password => 'mysql_root_password',
21        admin_email         => 'admin@iownz.you',
22        admin_password      => $admin_password,
23        keystone_db_password => 'keystone_db_password',
24        keystone_admin_token => $keystone_admin_token,
25        glance_db_password  => 'glance_pass',
26        glance_user_password => 'glance_pass',
27        nova_user_password  => 'nova_pass',
28        nova_user_password  => $icclab::params::nova_user_password,
29        rabbit_password     => $icclab::params::rabbit_password,
30        rabbit_user         => $icclab::params::rabbit_user,
31        export_resources    => false,
32
33    }
34
35    # Optional: include if you want authorisation information
36    #     stored in a local file, located in /root/
37    class { 'openstack::auth_file':
38
39        admin_password      => $admin_password,
40        keystone_admin_token => $keystone_admin_token,
41        controller_node     => $icclab::params::controller_node_internal,
42
43    }
44
45 }
```


What's in a compute node?

```
1 class icclab::compute{
2
3   include icclab::params
4
5   class { 'openstack::compute':
6
7     public_interface => $icclab::params::public_interface,
8     private_interface => $icclab::params::private_interface,
9     internal_address => $ipaddress_eth0,
10    libvirt_type      => 'qemu',
11    fixed_range      => $icclab::params::fixed_range,
12    network_manager  => $icclab::params::network_manager,
13    multi_host       => true,
14    sql_connection   => $icclab::params::sql_connection,
15    nova_user_password => $icclab::params::nova_user_password,
16    rabbit_host      => $icclab::params::controller_node_internal,
17    rabbit_password  => $icclab::params::rabbit_password,
18    rabbit_user      => $icclab::params::rabbit_user,
19    glance_api_servers => "${icclab::params::controller_node_internal}:9292",
20    vncproxy_host    => $icclab::params::controller_node_public,
21    vnc_enabled      => true,
22    verbose          => true,
23    manage_volumes  => true,
24    nova_volume      => 'nova-volumes'
25
26   }
27
28 }
```


Conclusions/Learnings

- Automation is essential
- Puppet codifies learnings, makes sharing easy
- Foreman a central management point, full lifecycle, adaptable to other services
- Dependence on infrastructure service management frameworks is lessened
 - Fast and efficient to install new ones with a tool chain
- Other than SLA guarantees, the only guarantee to maintain is the API between provider and customer and this is where standard APIs are need such as OCCI/CDMI/OVF.

Next Steps



- OpenStack to be rolled-out in ICCLab
 - New data centre, rolled-out within the month
 - Will include all OS Nova (Essex) and Swift services
 - Including OCCl interface
 - puppetlab-nova pull-request available



- OpenStack to be rolled-out in GWDG
 - Will include all OS Nova (Essex) and Swift services
 - Providing production-quality OpenStack services

Thanks!

Questions?

Everything Presented is Documented at:

<http://www.cloudcomp.ch>

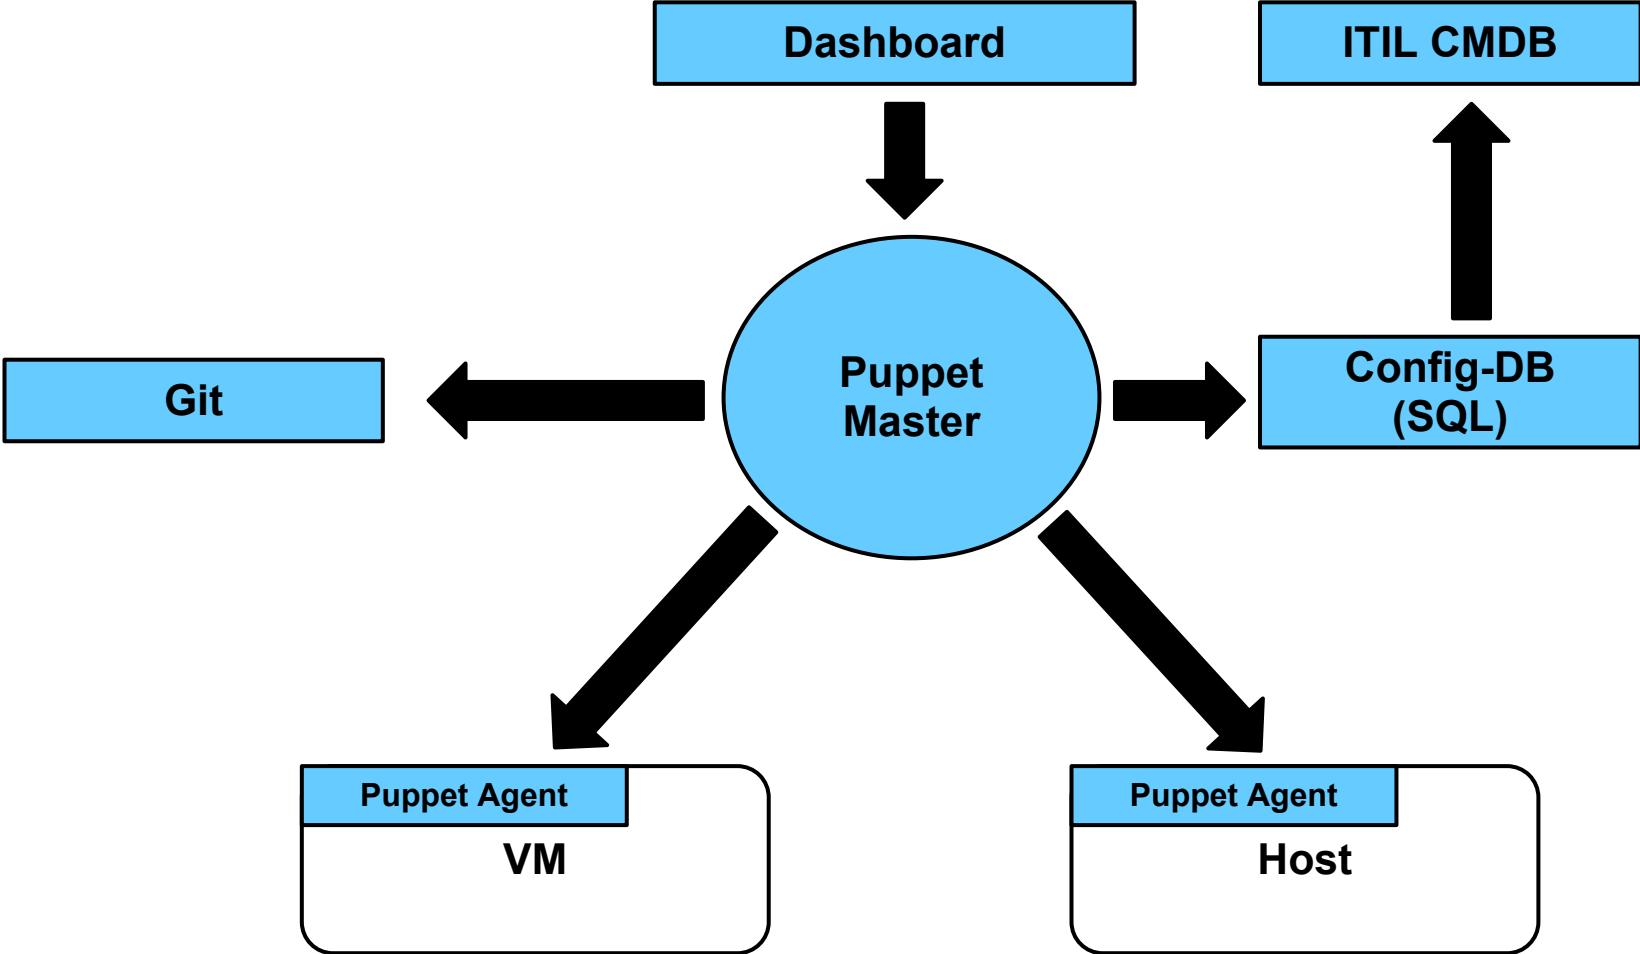
<http://cloud.gwdg.de>

Including:


- HOWTOs
- Foreman, Puppet, OpenStack installs
- Virtual Machine images

Backup slides

Toolchain map



Foreman Arch



fedora^f



Foreman

Netinstall (PXE)

TFTP

DHCP

DNS

HTTP

Bare-Metal

OS artefacts

kernel

packages

initrd

kickstart.ks

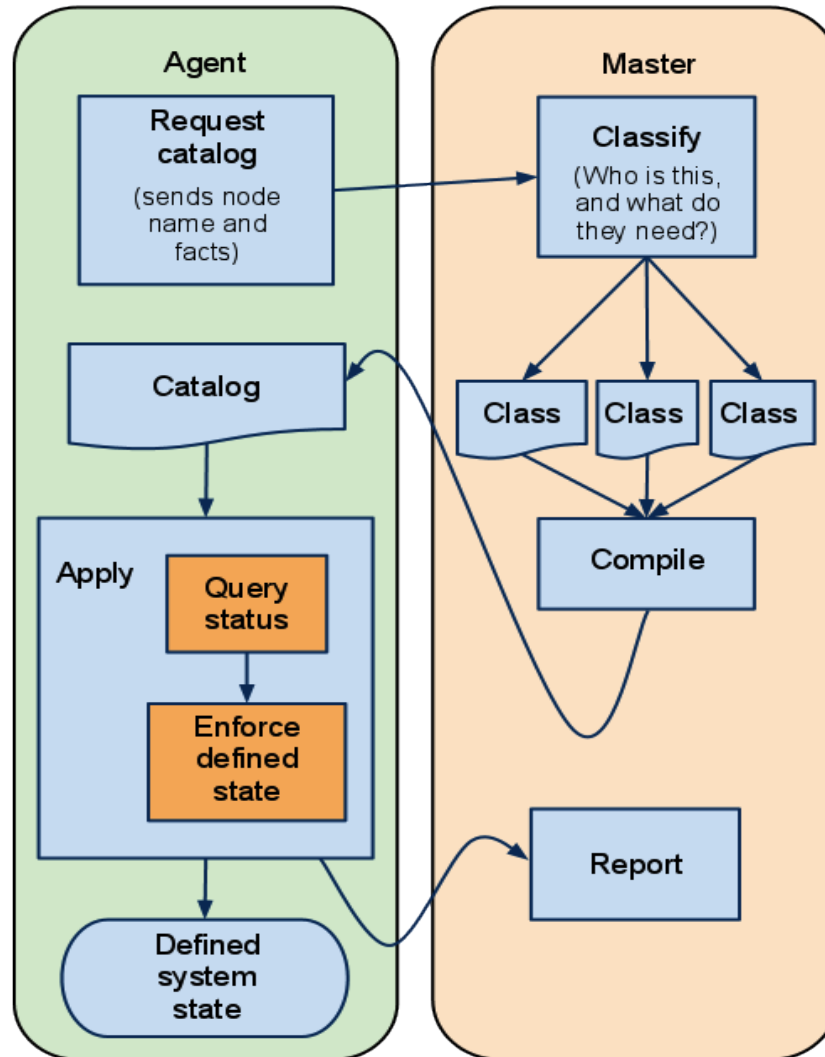
VM

XEN

VM

KVM

Puppetmaster <-> agent interaction



What are the common config params?

```
1 class icclab::params{
2
3
4 /* -----Shared Connection Settings-----*/
5
6 ##### Important to set! #####
7 $controller_node_address = '192.168.56.3'
8
9 $controller_node_public = controller_node_address
10 $controller_node_internal = controller_node_address
11 $sql_connection = "mysql://nova:${icclab::params::nova_db_password}@${controller_node_internal}/nova"
12
13 /* -----*/
14
15
16 /* -----Shared Auth Settings-----*/
17 $nova_user_password = 'nova_pass'
18 $rabbit_password = 'rabbit_pass'
19 $rabbit_user = 'rabbit_user'
20 /* -----*/
21
22
23 /* -----Shared Networking Settings-----*/
24 $network_manager = 'nova.network.manager.FlatDHCPManager'
25 $fixed_range = '10.0.0.0/24'
26 $public_interface = 'eth0'
27 $private_interface = 'eth1'
28 /* -----*/
29
30 }
```

GWDG Cloud topology

