Wikimedia Labs does things

- Vms ('instances')
- Projects ('tenants')
- User accounts (for Labs and elsewhere)
- Shared storage
- Database storage
- Database replicas
- Membership, auth, access among all of the above

Wikimedia Labs uses things

- Openstack (nova, keystone, glance, designate, horizon)
- MediaWiki (OpenStackManager)
- LDAP
- NFS
- MySQL

- **Today will mostly be about Openstack, VMs**

Create VMs

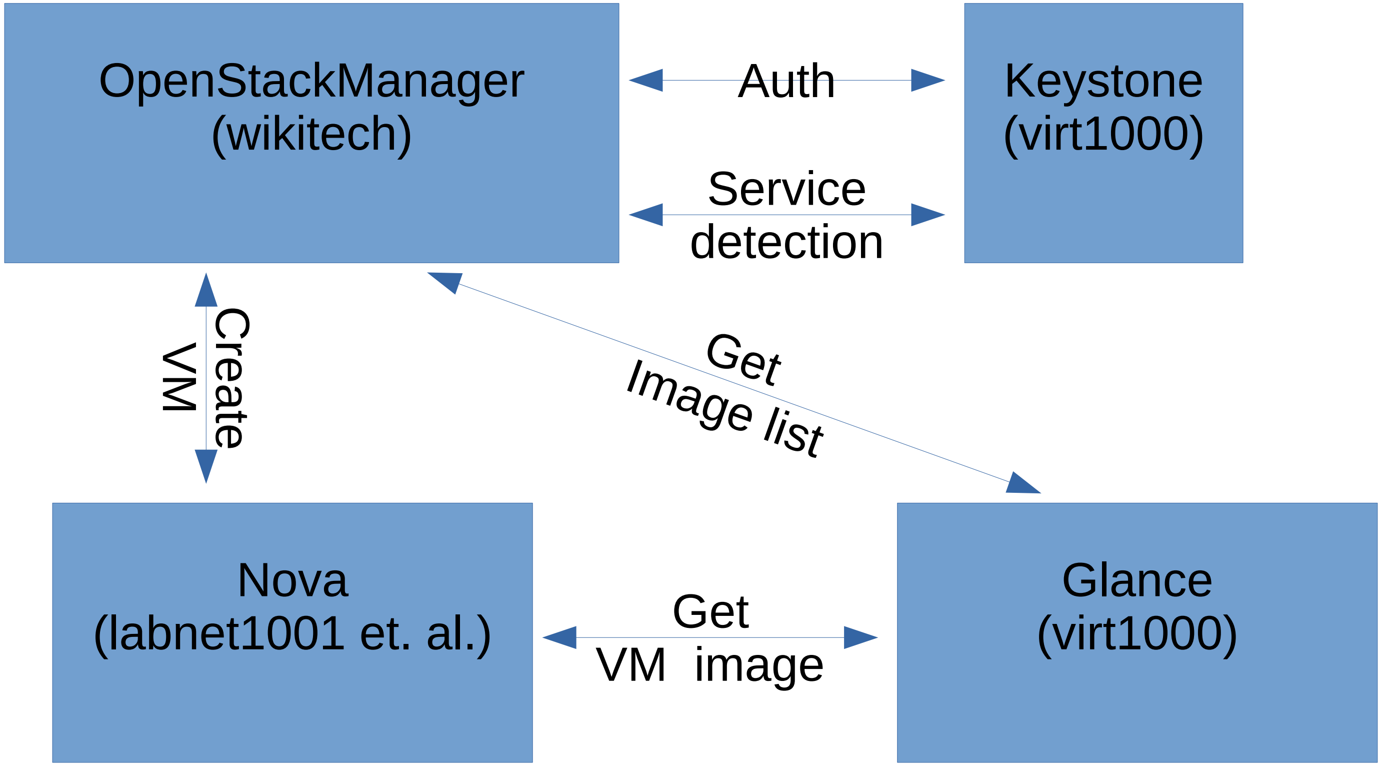

OpenStackManager (wikitech)

Keystone (virt1000)

Nova (labnet1001 et. al.)

Glance (virt1000)

OpenStack

- A giant set of Software-as-a-Service APIs and implementations
- Public-facing REST interfaces
- Internally, services communicate with RabbitMQ
- Written in Python
- Most services provide command-line front-ends to the REST apis. Such cli tools are all installed on virt1000.
- OpenStack provides APIs, interfaces, management. Generally not services themselves.

- Labs uses: Nova, Keystone, Glance
- Coming soon: Designate, Horizon

OpenStack Keystone

- Aka 'Identity', handles accounts and auth
- Users, Projects, Roles
- Service discovery

- Unintuitive but fairly simple
- Lives entirely on virt1000

OpenStack Glance

- File-store that Manages VM base images
- When you create a new image type (e.g. Jessie) you are adding it to glance.

- Very simple and stable
- Lives entirely on virt1000

OpenStack Nova

- Virtualization
- Many subservices: scheduler, conductor, api, network, compute

- Runs on many hosts: virt1000, virt10*, labnet1001
- All configured with /etc/nova/nova.conf, the same on all nova hosts

nova-scheduler

- Uses weighted algorithm to choose where to launch a new VM

- Runs on virt1000
- Very stable
- Start with the scheduler logfile when tracing a failed instance boot.
- /var/log/nova/nova-scheduler.log

nova-conductor

- Marshalls db calls from other services to the nova mysql db (also on virt1000)

- Runs on virt1000
- Very stable.
- /var/log/nova/nova-conductor.log
- I have never once looked at the conductor log.

nova-api

- Implements the nova REST interface and relays calls to other services via rabbit.
- Authenticates with Keystone

- Runs on labnet1001
- Stable
- /var/log/nova/nova-api.log
- A good place to look if wikitech or command-line commands get timeouts, 404s, 405s.

nova-compute

- Runs on virt1001-1012 (and soon, more)
- 'Hypervisor' wrapping KVM/qemu
- Creates, destroys, monitors VMs.
- Not involved in the active running of VMs. If it crashes, VMs are fine but nothing new can be scheduled on the hardware.

- Sometimes it locks up for no reason.
- /var/log/nova/nova-compute.log
- If you need to dive deeper into VM issues, dig into /var/log/libvirt

nova-network

- Configures bridge, dhcp server (and, for now, DNS server) for all VM network connections.
- Obsolete (now replaced by OpenStack Neutron) but it's still supported and works fine.

- Runs on labnet1001
- /var/log/nova/nova-api.log

Openstack Designate

- The future of Labs DNS
- Automatically manages private DNS records based on instance creation/deletion
- Backed with PDNS/mysql

- Everything is on Holmium, AKA labs-ns2

Openstack Horizon

- The future of Labs web UI
- Currently talks to Openstack Services
- Needs to write to ldap

- Everything is on Californium, AKA Horizon

Go ahead, restart it

- It is always safe to restart ANY of these openstack services.

- Nova is the api and configurator, not the actual VM service – that's handled by qemu, dnsmasq, etc.
- Rabbit and REST apis are both very tolerant of interruptions

…

- Next up: Account management

Auth apologies

- Labs predates Keystone and Horizon
- So all of this is a bit of a mess

- Mercifully, Keystone uses ldap as its store, so account information isn't duplicated

Manage accounts and privs

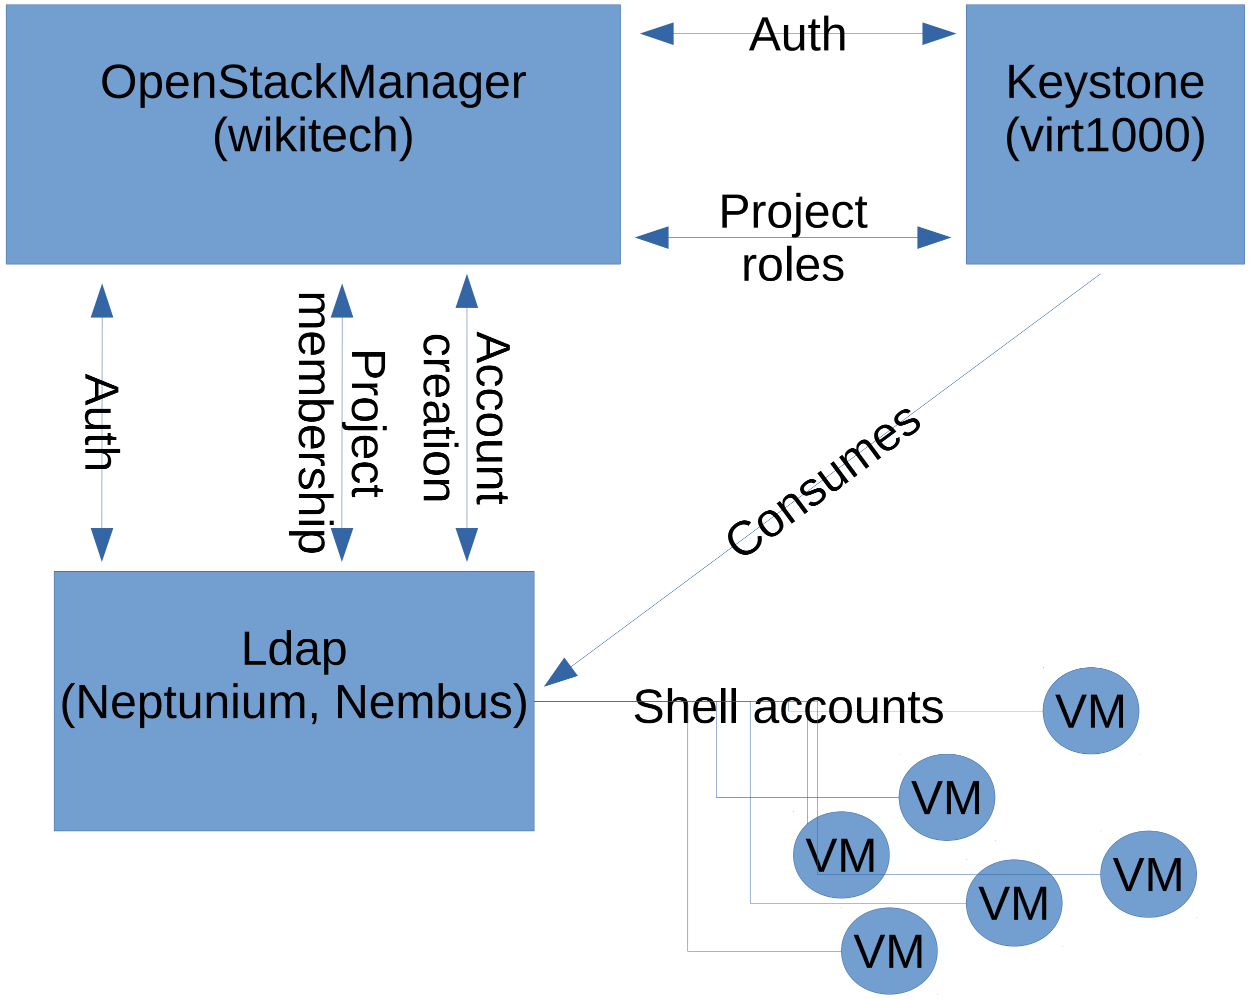

Keystone vs. Ldap

- Keystone is read-only, backed by ldap

  ...but...

- Wikitech login requires both keystone and ldap auth:
  - ldap auth for mediawiki actions
  - Keystone for OpenStack actions

- VM account access is just ldap

User accounts on VMs

- Accounts from pam_ldap
- Ssh keys from ldap (Trusty, Jessie)
- Ssh keys via shared mount (everywhere)

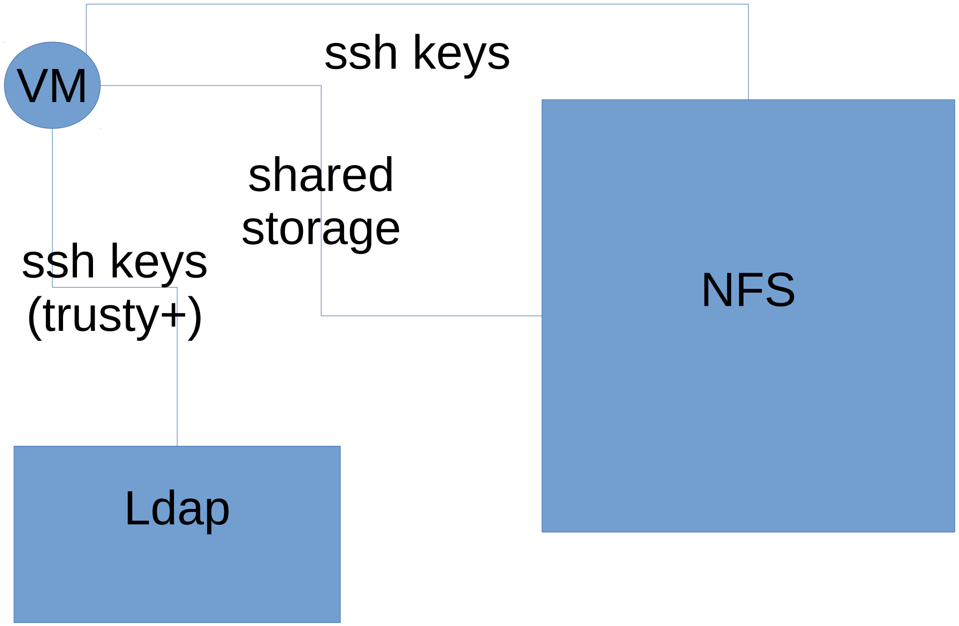

Keystone vs Wikitech

- Keystone enforces some access policies
- Wikitech enforces other access policies
- Sometimes they conflict

- This needs to be reconciled for Horizon

Day in the life (today)

- User asks to creates instance on Wikitech (thanks to ldap rights)
- Wikitech writes an ldap host record
- Wikitech writes the puppet node definition
- Wikitech creates a VM
- (hand-waving) Nova creates the VM
- Nova-network makes DNS entry in dnsmasq
- Nova updates the wiki instance page

Day in the life (future)

- User asks to create a VM on Horizon (thanks to keystone check)
- Horizon creates a VM
- Nova creates the VM
- Puppet node definition already exists, thanks to labs/project-scoped hiera
- Designate makes DNS entry in pdns
- Nova updates the wiki instance page

Topics for future sessions

- Toollabs
- NFS
- Databases