



redhat

Towards an Open Cloud Exchange

Kristi Nikolla, Mass Open Cloud
Monty Taylor, Red Hat




Northeastern




- Today's cloud landscape
- Use cases for an OCX
- The MOC
- Things we've done
- Lessons learned
- Possible next steps

Today's IaaS clouds

- One company responsible for implementing and operating the cloud
- Typically highly secretive about operational practices
- Exposes limited information to enable optimizations
- Vendor lock in: interfaces, price...



Is a different model possible? An “Open Cloud eXchange (OCX)”



Open Cloud Exchange as a solution

- An OCX provides a framework for solving these challenges:
 - Different participants in the cloud can stand up different services and make them available to the broader OCX userbase.
 - Multiple organizations can collaborate on a service offering, reducing operational cost while providing access to valuable data to the research and open source communities.
 - Open source software development can be informed by real user feedback!

Importance at even modest scale

- Enable cloud model for developing private cloud software & services
- Price is significantly cheaper
- Avoid vendor lock in
- Innovation & Research
 - Allow many providers to compete
 - Enable failure
 - Access to real data / real users / real scale
 - Ability to pilot systems to users

Why we care?

- Open Source Development and Research Requirements:
 - access to real data
 - access to real users
 - access to scale
- History tells us that when we open things up to rich communities it leads to competition and results in innovation and efficiency:
 - “The Cathedral and the Bazaar” by Eric Steven Raymond



CATHEDRAL

- Alibaba
- Amazon
- Baidu
- Google
- Microsoft

BAZAAR

- Mass Open Cloud
- NorthEast Storage Exchange
- FGPA Services
- GPU Services
- Open Datahub
- My cousins Minecraft Server
- Etc. etc.

Why Should the Community Care?

- The overall community has no place to run a Continuous Integration and Deployment of a cloud at scale so issues are not found until customers find them.
 - Issues because of interface changes that may be missed by different projects
 - Issues because they only show up at scale
 - Issues because the implementers were solving a different problem
- Quantity has a quality all its own
 - The problems change as usage grow

Use Cases

- Community CI/CD
- Infrastructure and Scale Labs
- Data Analytics and AI/ML
- Others?

Some Things We've Learned

- Many alternatives to public cloud offering exists, however it is difficult for operators to justify learning and operating a broad and diverse set of services.
- None of the clouds individually have the scales to attract the diversity of services of the proprietary clouds.
- Some customers want bare metal and containers - more want services they can just use
- **Without operating a cloud with real users, the open source community has found it challenging to address the user experience of real customers.**

About the MOC

- We are a nonprofit
- We want to build a public cloud that people can use, **and pay for**, so that:
 - We can pay our salaries (small team - 3 dev, 1 ops, 1.7 admin types, some students)
 - We can cover the costs of operating and maintaining/upgrading hardware as needed
 - We can do research on alternative economic models
 - **It will be easier for others to replicate (software based config)**
- All the above mean chargeback, showback and billing/reporting are required
- We want to build a public cloud with strong monitoring
- *Openstack was not built with these things in mind.*
- **We have a vision of an open cloud exchange, we want to make it real**

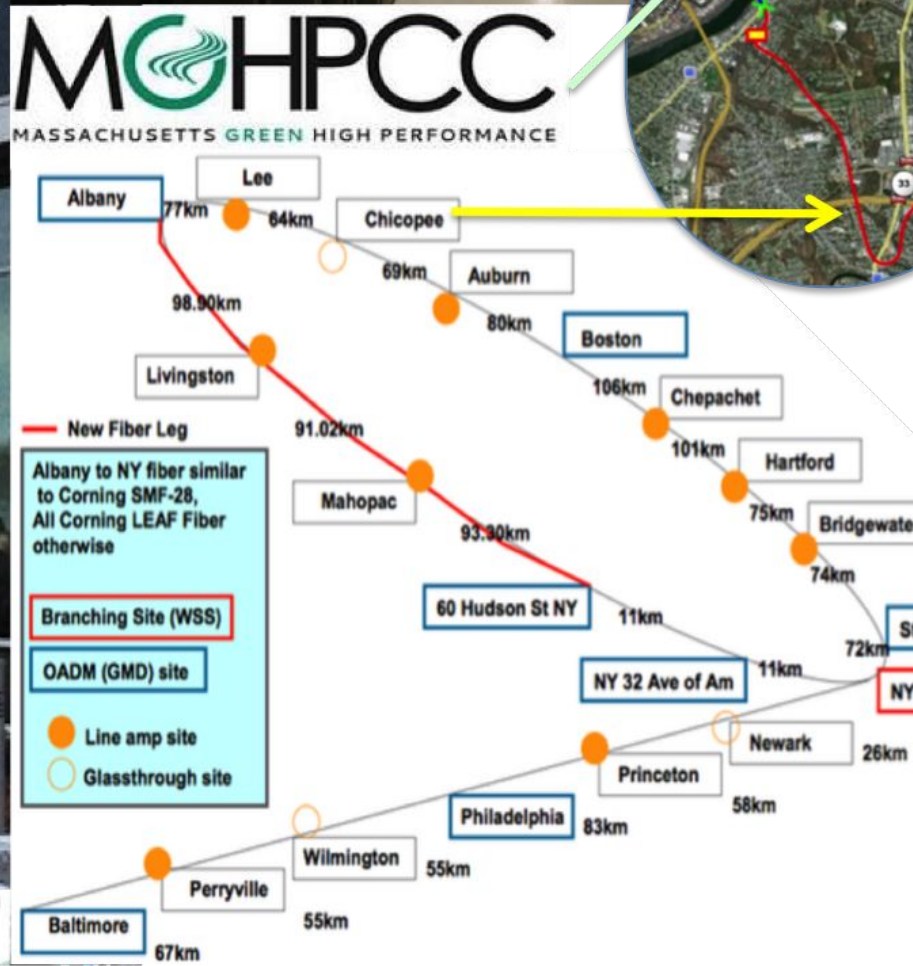

THE MASS Open Cloud COLLABORATORS

CANONICAL

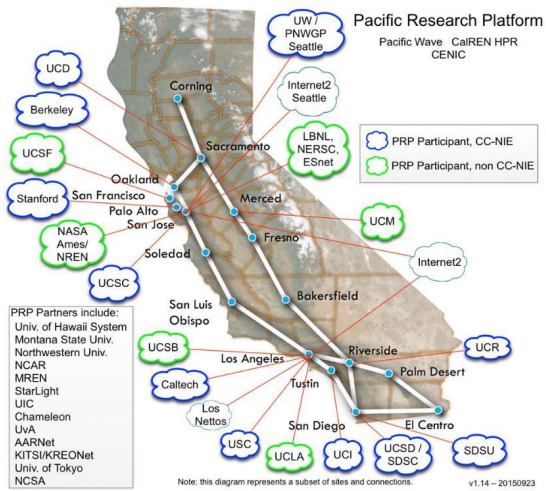


Northeastern





- 90,000 square feet + can grow
- 17+ thousand HPC users, potentially many more cloud users



Imagine shrinking Pacific Research Platform to the size of a building




Widths are proportional to enrollment

Consortia by Enrollment

MGHPCC Consortium comparable to Pacific Research Platform

- Huge community covering every field of research
- Collaborations across the globe
- Massive data and computational requirements
- Massive student population covering every discipline



Required capabilities for an MVP of OCX

1. An Elastic Secure Infrastructure for on-demand hardware use,
2. Production OpenStack, Ceph and Kubernetes services for both end users and higher level service offerings,
3. Single sign on (SSO) access to OCX services via public and private identity providers,
4. Resource federation between multiple OpenStack services,
5. A pricing guide and billing system.
6. A user management system

Required capabilities for an OCX

Once MVP is Done: Multiple Participants may:

1. Deploy their own hardware,
2. Deploy software services on top of hardware allocated dynamically for that purpose, and
3. Charge for these services.

Required capabilities for an MVP of OCX

1. **An Elastic Secure Infrastructure for on-demand hardware use,**
2. Production OpenStack, Ceph and Kubernetes services for both end users and higher level service offerings,
3. Single sign on (SSO) access to OCX services via public and private identity providers,
4. Resource federation between multiple OpenStack services,
5. A pricing guide and billing system.
6. A user management system

Secure Elastic Hardware

Solution for securely bursting nodes between different clusters.

Open Source Components:

- Network booting and provisioning (M2)
 - HIL - Network Isolation
 - M2 - Network booting and image management
- Bolted - isolated enclave of physical machine for security
 - Combines HIL, M2 and Keylime attestation from MIT Lincoln Labs

Required capabilities for an MVP of OCX

1. An Elastic Secure Infrastructure for on-demand hardware use,
2. **Production OpenStack, Ceph and Kubernetes services for both end users and higher level service offerings,**
3. Single sign on (SSO) access to OCX services via public and private identity providers,
4. Resource federation between multiple OpenStack services,
5. A pricing guide and billing system.
6. A user management system

Production

- Production OpenStack Deployment - Kaize
 - Had 576 cores
 - An additional 1620 cores will soon be available
- Started at Kilo and upgrade to Pike
 - Planned redeployment with Rocky. Deploying OpenStack has gotten a lot easier.

Production

- Pure Research OpenStack Cloud - Engage1
- Production OpenShift running on OpenStack
 - Upgrades have been hit or miss

Production: Lessons Learned

- *Velo-ceph*-raptors
- User support is time consuming

Required capabilities for an MVP of OCX


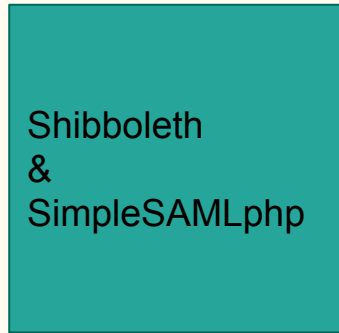
1. An Elastic Secure Infrastructure for on-demand hardware use,
2. Production OpenStack, Ceph and Kubernetes services for both end users and higher level service offerings,
3. **Single sign on (SSO) access to OCX services via public and private identity providers,**
4. Resource federation between multiple OpenStack services,
5. A pricing guide and billing system.
6. A user management system

Single-Sign On

Allow users and researchers to use their institutional account to log in.
For universities we federate with InCommon.



Single-Sign On




- Things are more complicated than they look

Required capabilities for an MVP of OCX

1. An Elastic Secure Infrastructure for on-demand hardware use,
2. Production OpenStack, Ceph and Kubernetes services for both end users and higher level service offerings,
3. Single sign on (SSO) access to OCX services via public and private identity providers,
4. **Resource federation between multiple OpenStack services,**
5. A pricing guide and billing system.
6. A user management system

Resource Federation with OpenStack

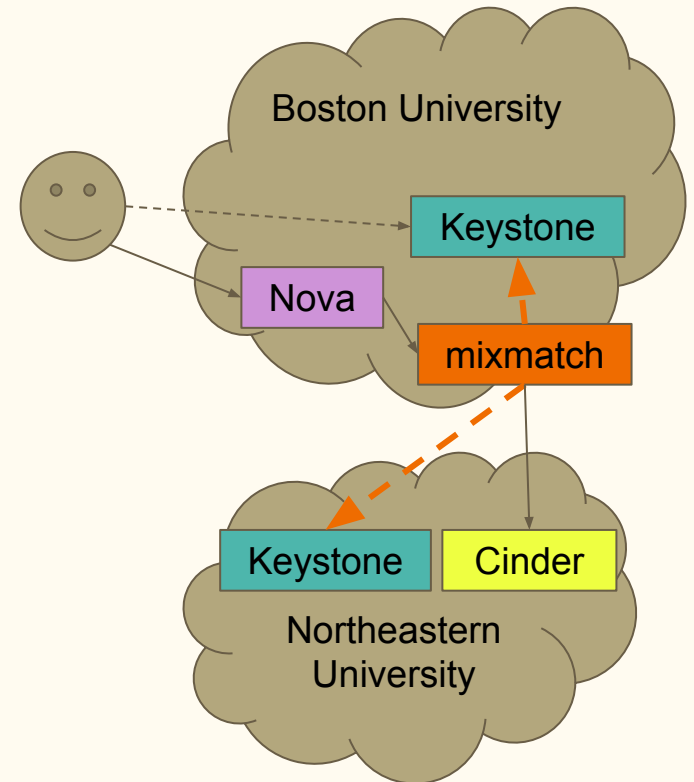
Enables using resources (such as Cinder volumes, Glance images, etc) to be shared across OpenStack deployments.



Resource Federation with OpenStack

Open Source Component:

- Mix & Match



Required capabilities for an MVP of OCX

1. An Elastic Secure Infrastructure for on-demand hardware use,
2. Production OpenStack, Ceph and Kubernetes services for both end users and higher level service offerings,
3. Single sign on (SSO) access to OCX services via public and private identity providers,
4. Resource federation between multiple OpenStack services,
- 5. A pricing guide and billing system.**
6. A user management system

Pricing

Pricing - what we learned - it's complicated :-)

1. We expect to be able to charge $\sim\frac{1}{2}$ to $\frac{1}{3}$ of commercial services,
2. Since research based we need to figure out models where researchers may purchase hardware since some grants favor that,
3. Competing on price is less interesting to our users than providing services which let them solve a problem. Most researchers and small businesses want to solve problems not focus on networking, etc.

Project Summary Bill

Goal is Single Page per project							
Project Name	Sponsoring Organization	Project Lead Name	Project Lead Email				
Per VM	MVP : 1 page per project Project Name Sponsoring organization Project Lead Name Project Lead Email Per VM Compute Usage Per VM Memory Usage Totals for these			memory usage	Network Usage	<i>Note - some of these will not be available initially</i>	
VM1				M1 memory	Network Usage	<i>Note - some of these will not be available initially</i>	
VM2					Network Usage	<i>Note - some of these will not be available initially</i>	
Total VM Usage							
Per Container				container memory	network usage		
Obj1							
Total Object Storage							
Total Volume Storage							
Persistent Volumes (Container Volume)							
persistent volume 1							
persistent volume 2							
Floating IP addresses							
OpenShift Usage (pointer to Placeholder Tab which							
Quota from Jan1-Jan15				object storage	network	memory	
Quota from Jan15-Jan30				object storage	network	memory	
						block storage	
						block storage	
List of Users							
Total Usage	Compute	Block Storage	Object Storage	Network Usage	memory		
Cost						<i>Note-needs to be figured out based on usage once we know it</i>	

Projects by Institution

Projects by Institution	Project Lead Name	Project Lead Email	Compute	Block Storage	Object Storage	Network Usage	memory	quota charges			
Project Name 1	Project Lead Name	Project Lead Email	Compute	Block Storage	Object Storage	Network Usage	memory	quota charges			
Project Cost	--	--						<i>Note-needs to be figured out based on usage once we know it</i>			
...											
Project Name N	Project Lead Name	Project Lead Email	Compute	Block Storage	Object Storage	Network Usage	memory	quota charges			
Project Cost	--	--						<i>Note-needs to be figured out based on usage once we know it</i>			
Total Usage	Compute	Block Storage	Object Storage	Network Usage							
Cost								<i>Note-needs to be figured out based on usage once we know it</i>			

MVP:

Project Name

Project Lead Name

Compute Usage

Memory Usage

We have succeeded in proving that those 10,000 ways will not work.

–Apologies to Edison

Telemetry Data

- As we discussed earlier - clouds are locked down and limited information is available for optimization and research
- As an open cloud we want that information to be available
- So we set about providing it
 - Ceilometer (2015-2016)
 - VERY SLOW
 - It made our network unusable
 - After Ceilometer we investigated Monasca
 - It was OK, but when we came back around to billing it was a non-starter

Seems pretty basic

- We planned to pull the data together ourselves and generate the reports (by now it was late 2017)

Area where we will need the most help!

Required capabilities for an MVP of OCX

1. An Elastic Secure Infrastructure for on-demand hardware use,
2. Production OpenStack, Ceph and Kubernetes services for both end users and higher level service offerings,
3. Single sign on (SSO) access to OCX services via public and private identity providers,
4. Resource federation between multiple OpenStack services,
5. A pricing guide and billing system.
6. **A user management system**

User Onboarding and Signup

Simple Use Case:

As a new user I want to be able to Sign onto the MassOpen Cloud and sign up for multiple services (storage, computer, partner services).

As a user leaves the MassOpen Cloud those resources should automatically be returned for other users.

To Support Users from Multiple Institutions and Companies - User Management Needed

- Need to be able to
 - Keep Track of Users across services for Billing and Permissions
 - Chargeback and Billing
 - *By User*
 - By Project
 - By Organization
 - By Resource
 - **Users of Projects and Organizations should be able to *sign up for* resources without requiring human intervention**
 - Monitoring and Auditing
- Support Users

Once Onboard

- Need to be able to
 - Keep Track of Users across services for Billing and Permissions
 - **Users of Projects and Organizations should be able to *sign up for resources without requiring human intervention***
 - Release Resources associated with Projects/Users
 - Monitoring and Auditing
 - Support Users

What else have we been doing?

- Integrate FPGA's and GPU's into OpenStack and OpenShift.
- Interviews with Startups and Small Manufacturers
- Outreach to Research IT organizations

More things we've learned

- Any time we think a project or RHEL feature will work we ask “are any current customers using it this way”
 - Of course we don't always wait for the answer
- If our Senior Infrastructure Engineer is out of town, bad things will happen.
- The community has been incredibly supportive and giving - thanks!

We Want Your Help!

<https://massopen.cloud>