Smart. Safe. Together

MIT MANAGEMENT SLOAN SCHOOL

IT Partners Conference :

Cloud Computing for Research

June 14th, 2022

Rajiv Shridhar

Senior Director for Strategic Technologies &
Information Security Officer

CONNECTING TO PUBLIC CLOUD FROM ON-PREMISE COMPUTING RESOURCES

- Burst to the cloud for “elastic” HPC capacity
- Wide-ranging non-HPC compute services
- Accessing public datasets
- Archival storage

- Of course, one can navigate the web management consoles of these cloud services. And use the cloud accounts that Rikus demonstrated for consolidated / discounted billing.

- For automation, particularly for research computing, we’d like to have other options than the web interface.
- And ways to run HPC clusters on the public cloud.

COMMAND LINE ACCESS

Amazon, Microsoft and Google all provide CLIs.

Supports the functionality of the public cloud management console using remote terminal

- Linux, MacOS and Windows versions
- Typically integrates with Python and/or Windows PowerShell
- Easy to install and automate computing processes
- Extend the on-premise Engaging cluster into the public cloud environment:
  - *For "burst to the cloud" computing; transferring data to/from public cloud vendor storage; provisioning/deprovisioning, managing and updating cloud compute, network, databases, security and other resources*

e.g.
```
# aws ec2 create-image --instance-id i-44a44ac3 --name "DevSloanRC AMI" --description "AMI for dev FSA server"          % (create a new image)
# aws ec2 start-instances --instance-ids i-dddddd70                                                                    % (launch a new EC2 instance)
```

But what if I wanted to use the web/GUI for access and work? But I want to use it for research computing (e.g. HPC clusters)

Also, anything beyond basic cloud usage presents a non-trivial learning curve!

High Technical Barrier to Entry, especially for new users!

RONIN cloud service experimentation at MIT

RONIN (https://ronin.cloud) is a service that provides simplified layer on top of AWS EC2, VPC, Route 53, IAM targeted to common research computing needs.

RONIN is designed to lower the technical barrier to entry

Features

Integrated MIT auth sign-in to service

Simplified resource interface, cost reporting and monitoring

Super easy start, stop, restart VMs

Desktop app for GUI interaction

Elastic Slurm clusters (with spot pricing options)

… experimental – integrated automated https endpoint, integrated VM creator MIT auth, “under development”

https://ronin.mitresearch.cloud/login.php

Courtesy Chris Hill, cnh@mit.edu

RONIN - integrated MIT auth sign-in to service

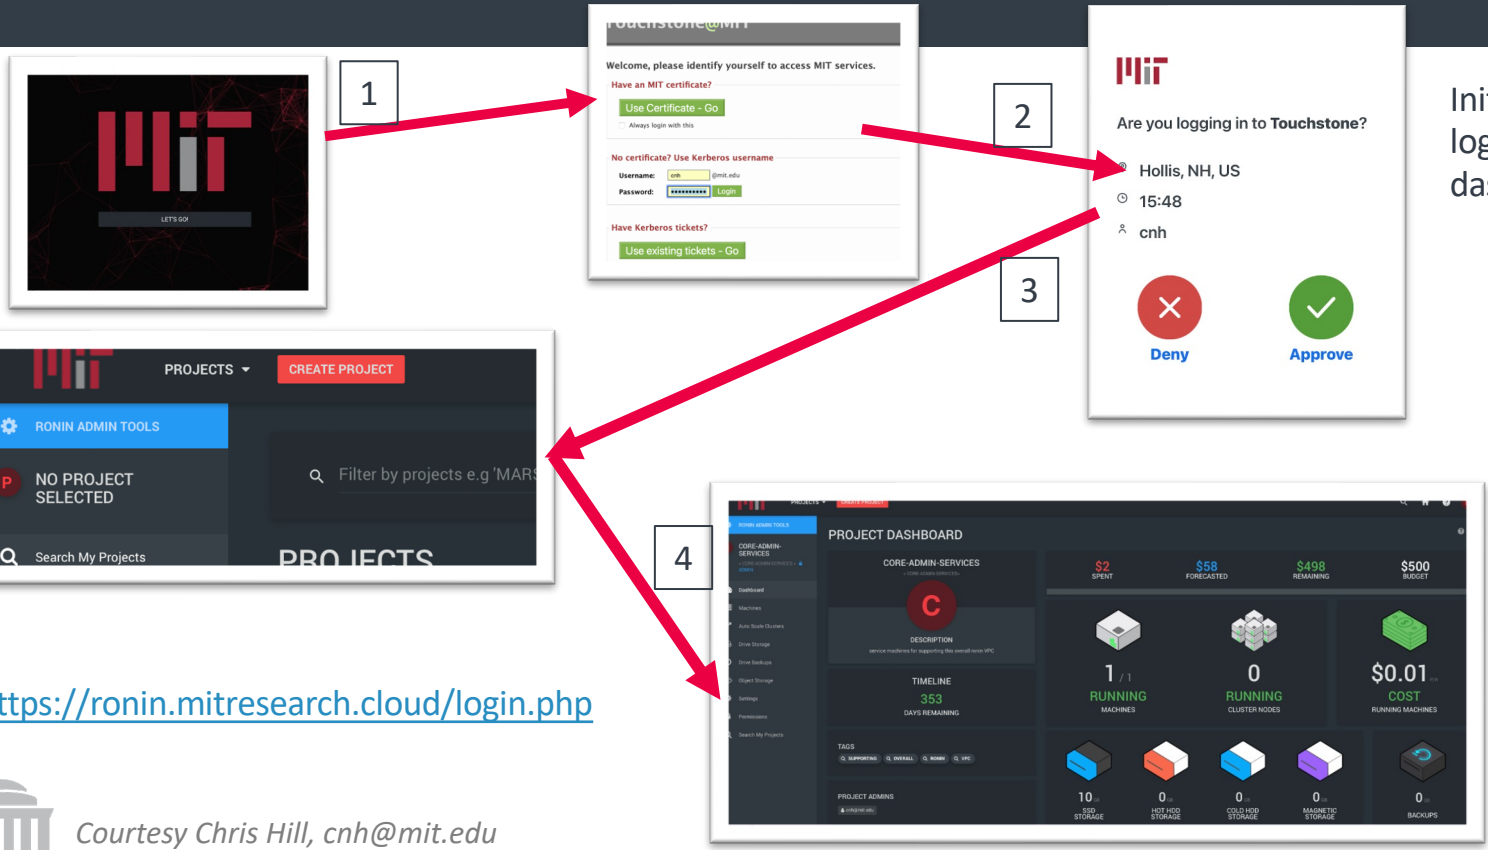

Initial steps to log in to RONIN dashboard

https://ronin.mitresearch.cloud/login.php

RONIN - Simplified resource interface, cost reporting and monitoring

RONIN organizes resources around projects

a project has a budget, budget alerts and cost summaries

RONIN simplifies AWS EC2 to a minimal set of resources e.g. machines, storage, clusters, images and presents these on visual dashboards.

More advanced AWS features are not accessible

Runs in an AWS account behind the scenes.

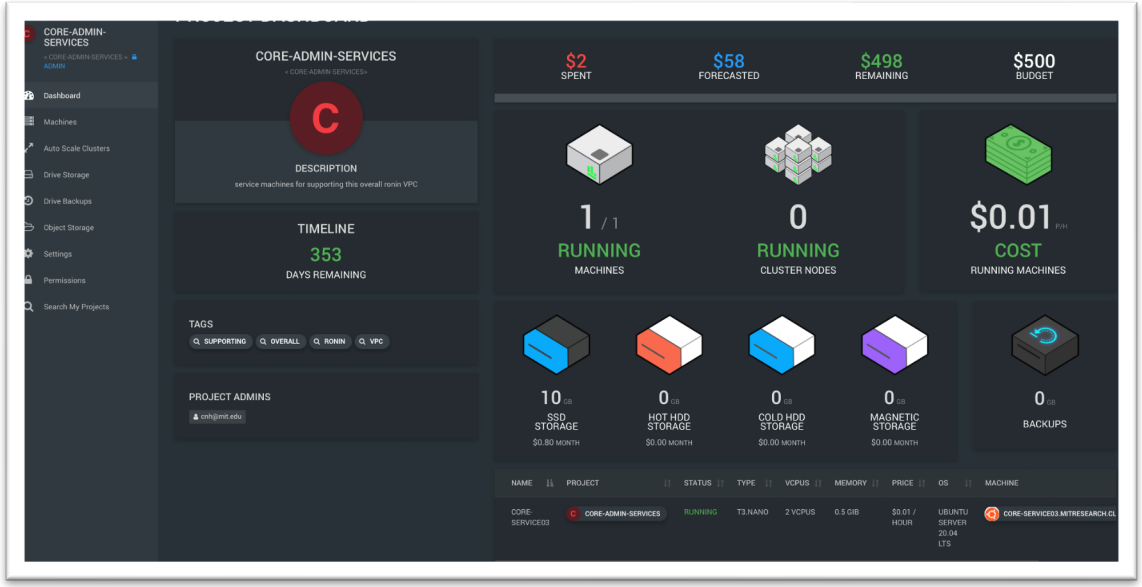

RONIN – super easy to start, stop, restart VM

RONIN VMs launch with DNS all set

Can ssh in and configure as needed

Simple interface to start, stop and restart. So can switch off when not is use – useful for expensive resources (for example on-demand A100 GPU system).

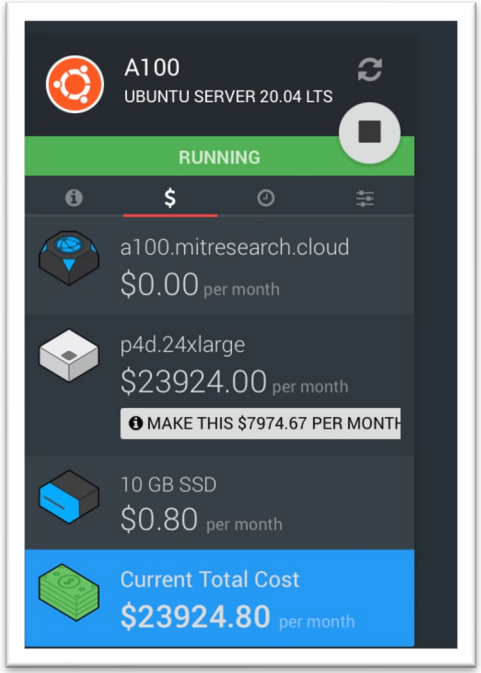

Cost per month running $23,924.80/month

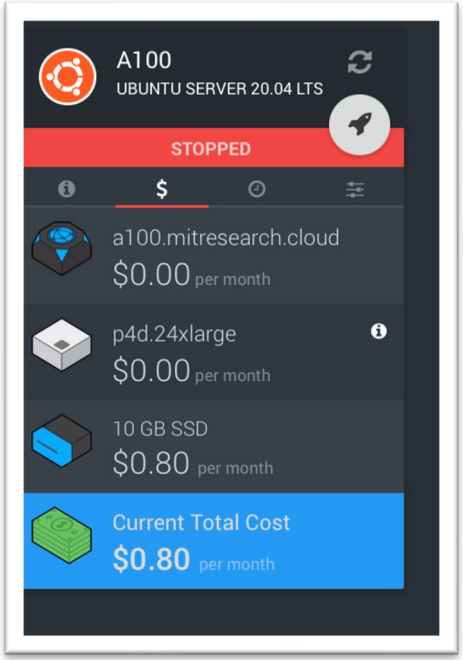

Cost per month stopped (but ready for restart) $0.80/month

RONIN – other features

Automated stop-restart on a schedule (e.g. expensive on-demand GPU machine on for Tue,Thu class, 12-4 only ➔ $1,200/month v. $23,924/month).

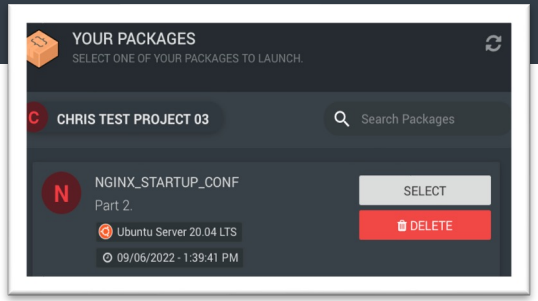

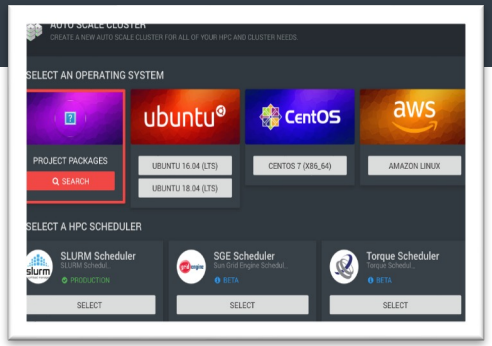

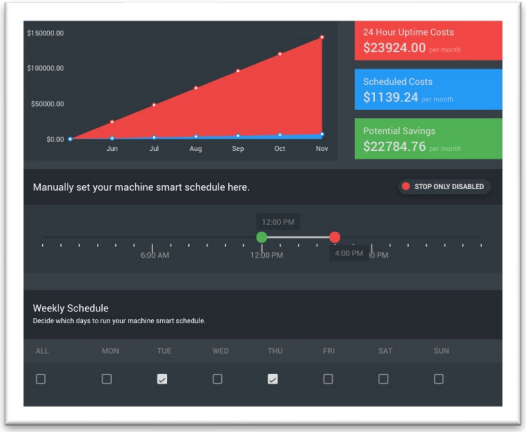

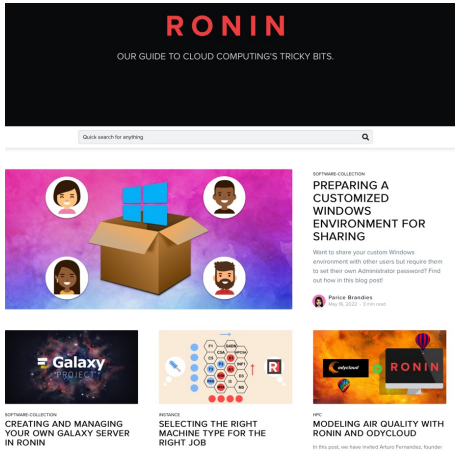

Elastic Slurm clusters. Compute nodes launch when there is demand (can use spot price bid).

Friendly documentation https://blog.ronin.cloud

RONIN (and AWS) pricing

- **Inherits same pricing discounts at other MIT AWS**
- **Typical rule of thumb is spot-price machines used less than 25% of time/capacity are lower than 5year purchase.**
- **Very cost effective for occasional individual special system needs (large memory, many GPUs, lots of processing for just a few days etc…)**
- **For intense day-in, day-out use with fairly static resource needs on-premise cloud-like systems (Engaging, Satori, Supercloud) significantly lower cost.**
- **RONIN fee TBD, est $200/user/year**

Cost example (not considering F&A)

Five year cost of 4 x A100@80GB server v cloud est. spot-price equivalent

Avg utilization (24x7, 365)	Cloud Cost[1]	Engaging shared cluster total cost[2]	Difference
16%, 4 hours/day	$44,862	$61,376	Cloud saving of $16,513
25%, 6 hours/day	$67,293	$61,376	On-prem saving of $5.917
80%, 19.2 hours/day	$215,338	$61,376	On-prem saving of $153,962

[1] does not include services F&A, assumes spot-price is ~= reserved instance price (which historically is close).

[2] total cost i.e. capital purchase + 5y warranty, est. pro-rated space, power, networking, base systems support and operations.

Courtesy Chris Hill, cnh@mit.edu

RONIN – next steps

Currently single region, single AZ

If resources unavailable in AZ (e.g. 8 A100 GPU server) then no easy way to switch region/AZ.

Currently single AWS charging account

RONIN is set up on an AWS charging account, creates some management resources

No mechanism currently to have multiple RONIN instances for different AWS-MIT accounts

Looking for small number of pilot test projects starting August/September

Ideally local support as well as end-use tester

Able to provide cost-object for recovering AWS charges

A couple of other RC projects of interest

More compute ...

More data ...

top500

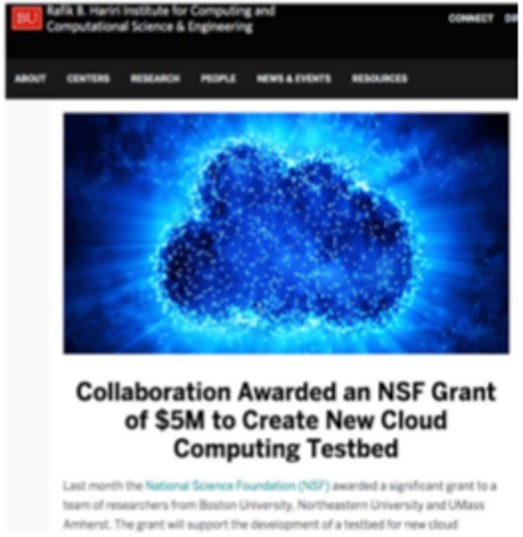

Northeast Storage Exchange

https://nese.mghpcc.org/

Based on an NSF DIBBs project (MGHPCC and member universities)

Ceph-based solution to drive:

1. Capacity : Towards Exascale, Largest known is CERN 100PB
2. Cost Effective : Open Source Software, unified storage management (object, block, file), No Single Point of Failure, can handle mixture of hardware in pools. *(Costs as low as $25/TB/year across 5 years)*
3. Secure Access : key based access either CephX or S3, dm-crypt of all drives, CephFS directories can pin to pools to separate client access.
4. Bandwidth : performance improves with increases in storage devices, load balanced,
5. Archival: a place to store data for long-term (using TAPE)

NESE: Some Use Cases

Support S3 Object Stores (default)

- CEPH S3

File-Based Access

– using CephFS + Globus

Replace Enterprise SAN/NAS

– using CephFS + RBD

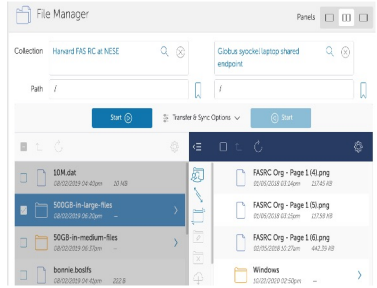

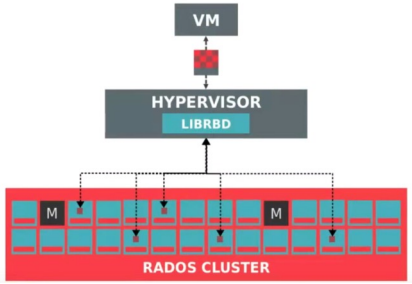

NESE includes federated identity management, billing, monitoring and operational support

Archival Storage: NESE Tape

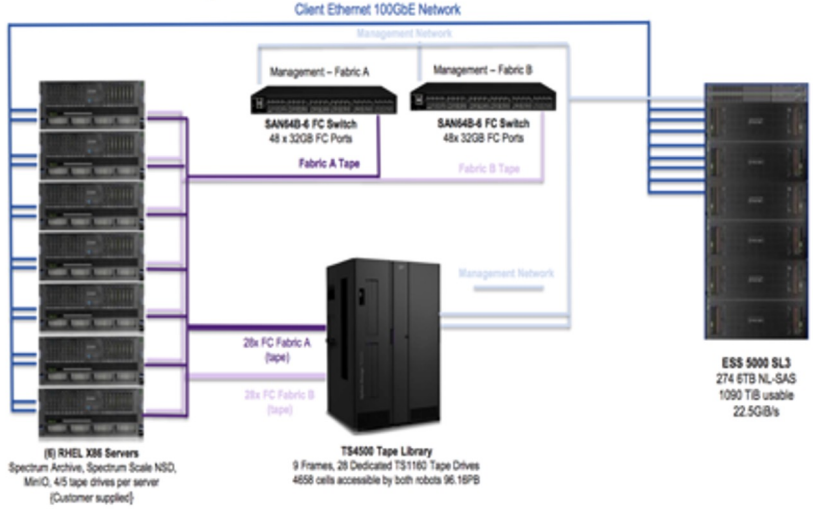

Nine frame IBM TS4500
Library + ESS-5000 front end

133 PB capacity
Max I/O: 11GB/s, scale up with drives.

Front end cache: 1 PB useable.

Protocols: POSIX and S3 via MinIO.

Software licenses and all maintenance paid for 5 years.

Extremely advantageous media buy-in price.

Immediate 50 PB buy-in from ATLAS.

Row 7, Pod A
Space for four libraries
Close to UPS power
Away from "transit" cable paths

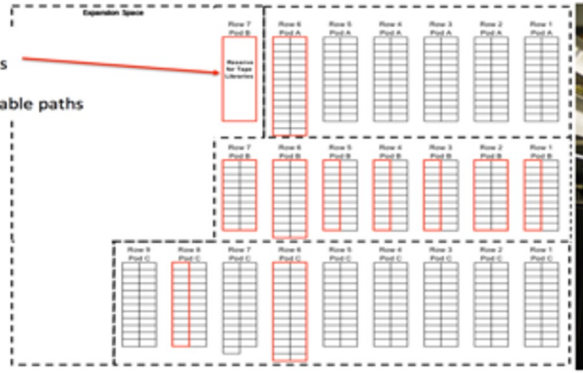

New England Research Cloud

Open source is key for research

- Most of Research IT depends on open source software and many of the major open source projects today emerged from the research and HPC world
- As Research IT moves to a cloud model, open source becomes even more ubiquitous and important
- Research projects in cloud now find their way into practical use through the efforts of researchers and industry. Ceph is a great example, but so many others (Elastic Secure Infrastructure at MOC, Chameleon at UIUC, etc.)
- Some research cannot be conducted in the commercial cloud instances for practical reasons, e.g. deep packet inspection for network traffic between applications.

NERC Goals

Build a cost-effective professionally operated on-prem cloud service for researchers that provides:

- **Self-service SaaS for easy access,**
- **Automated PaaS for custom workflows**
- **And IaaS that includes emerging technologies for hardware acceleration**

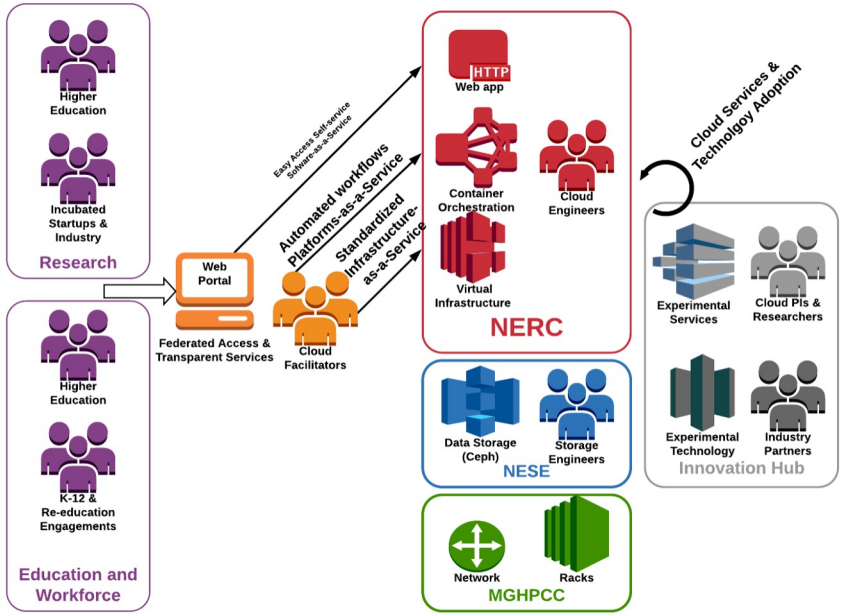

Under development with Harvard and BU as lead

Collaboration with many other open-source RC projects and vendors

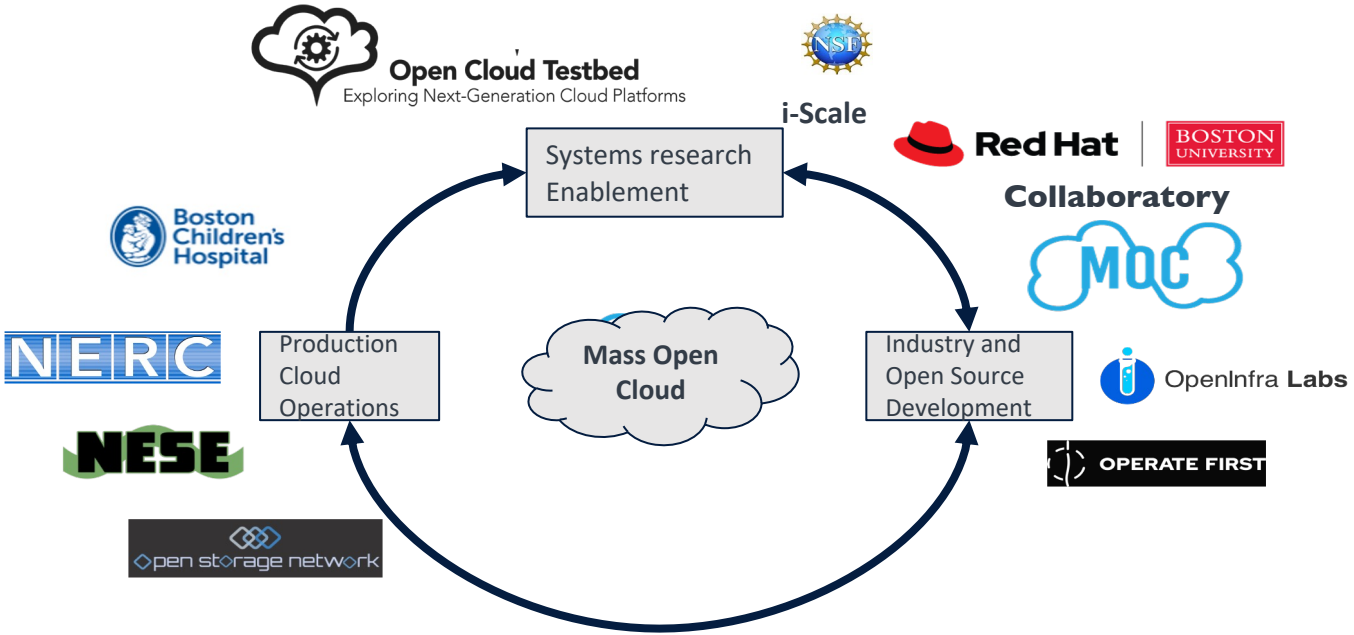

NERC represents OpenStack (transitioning to OpenShift) computing environments

QUESTIONS?