Adapting to the Evolving Research Landscape: Lessons Learned from Deploying OpenStack and Navigating the Challenges of Infrastructure as a Service

Benjamin Lynch
Graham Allan
James Wilgenbusch

> 900 Research Groups

119 Departments

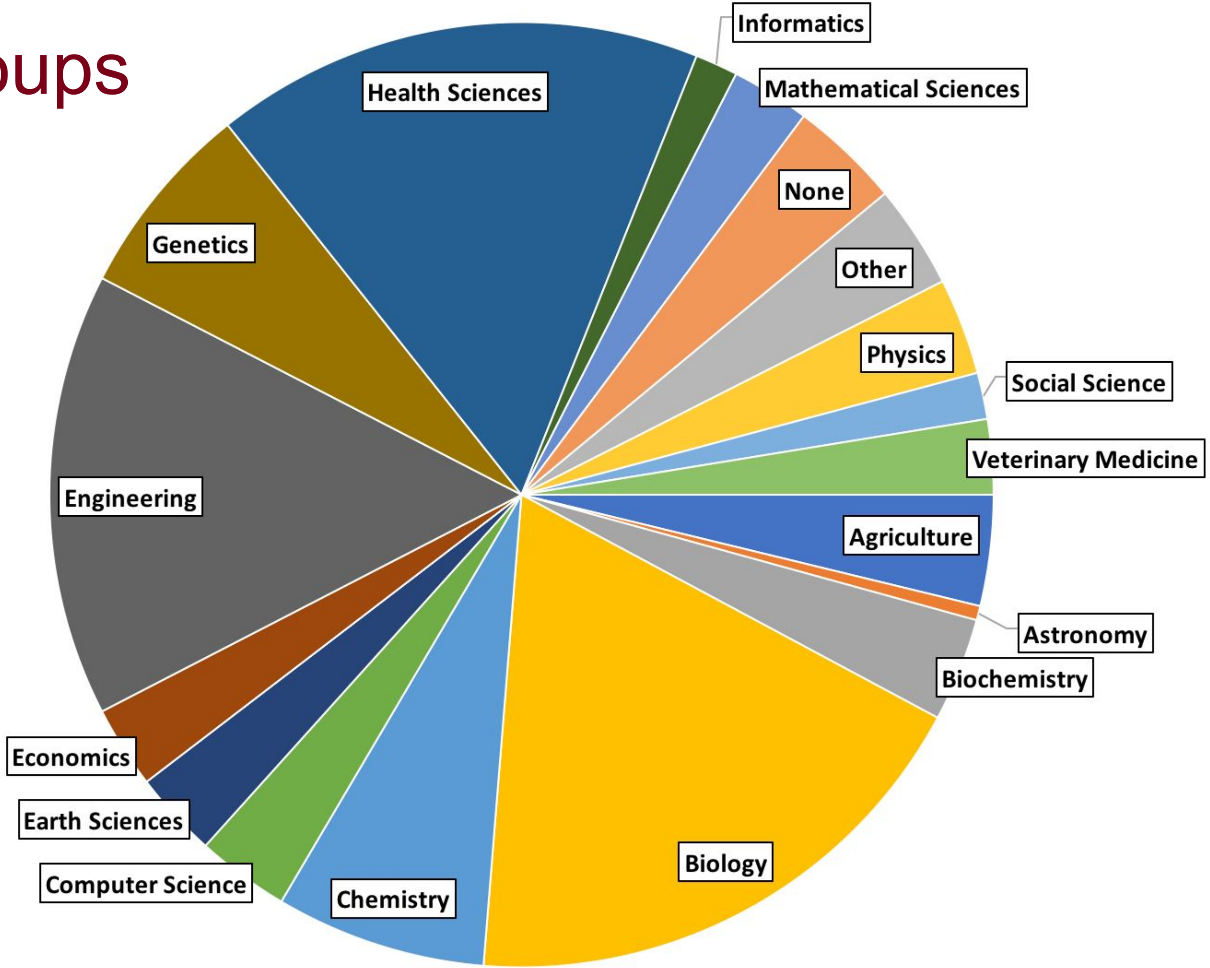

Supporting the Research Computing Needs of Researchers

- Determine the resources needed to solve the problem
- If we don’t have the resource, recommend someone that does

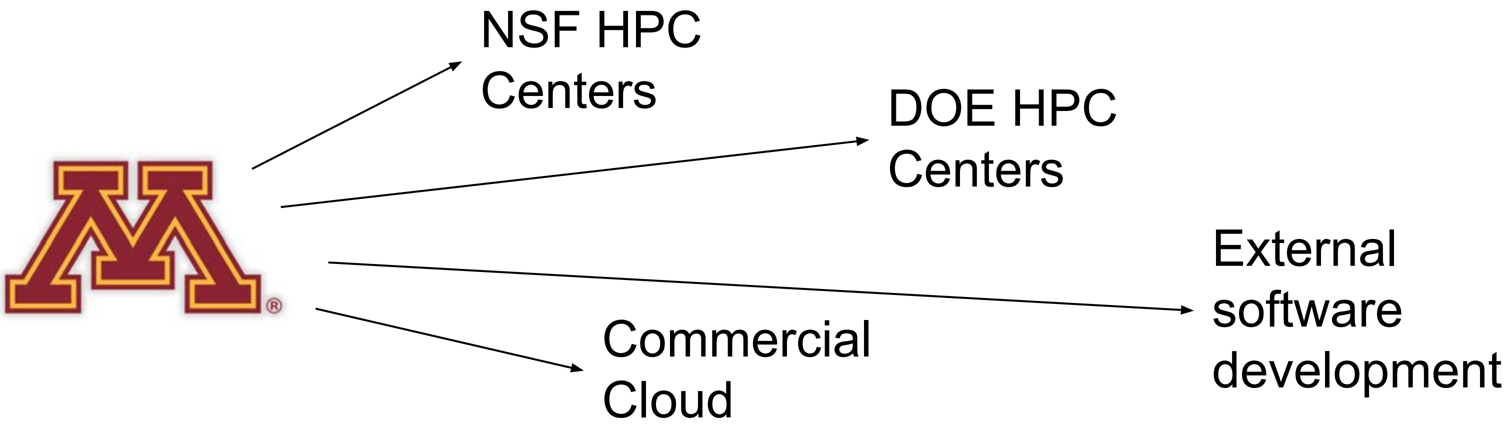

Supporting the Research Computing Needs of Researchers

- Determine the resources needed to solve the problem
- If we don’t have the resource, recommend someone that does

- If we don’t have a recommended solution, the researchers **WILL** find a way. It may include:
    - Unmanaged local & remote hardware
    - Inappropriate use of available resources

Identified Needs

- Bespoke data compliance and other needs that were not met in the existing HPC cluster
- Researcher-controlled persistent services and long-running jobs
- Network isolation, clinical pipelines
- MSI DevOps-maintained applications
  - Public gateways for global collaboration

Timeline

• OpenStack Experimentation	2011
• Internal OpenStack	2013

Local Cloud Computing

Goals

- Offer a lower cost than commercial cloud
- Reduce data movement
- More control over data
- Recover cost

Not Goals

- Scaling to 1000s of VMs
- Expanding free services

Why OpenStack?

Price

- Free and OpenSource
- Resellers can support enterprise deployments

Maturity

- NASA and Rackspace started the OpenStack project in 2010

Worldwide Popularity

- NASA, CERN, Bionimbus PDC (NIH), JetStream (NSF)
- BestBuy, Target, WalMart, Dreamhost, many more

Versatility

- OpenStack can run VMs on top most commodity and enterprise hardware
- Reuse infrastructure building blocks like existing HPC nodes

Timeline

- OpenStack Experimentation 2011
- Internal OpenStack 2013
- IaaS for researchers 2017

Stratus

- 7x Control nodes
- 20x Compute Nodes
  - 5 TB RAM
  - 560 Cores
  - Over-subscription rates: 4x CPU; 1.4x RAM
- 200 TB Ceph Block Device Storage
  - Booted VMs, Ephemeral and Persistent Data Volumes, Raw Images
- 512 TB Ceph Object Storage
  - S3 Interface
- 2x 40 GbE Network Switches

Stratus-Dev

- 2x Control nodes
- 2x Compute Nodes
  - 44GB RAM
  - 16 Cores
- 91 TB FrankenCeph Block Device Storage
- 1 GbE Networking

Local Cloud Deployment

- OpenStack - Newton
- Ceph Storage

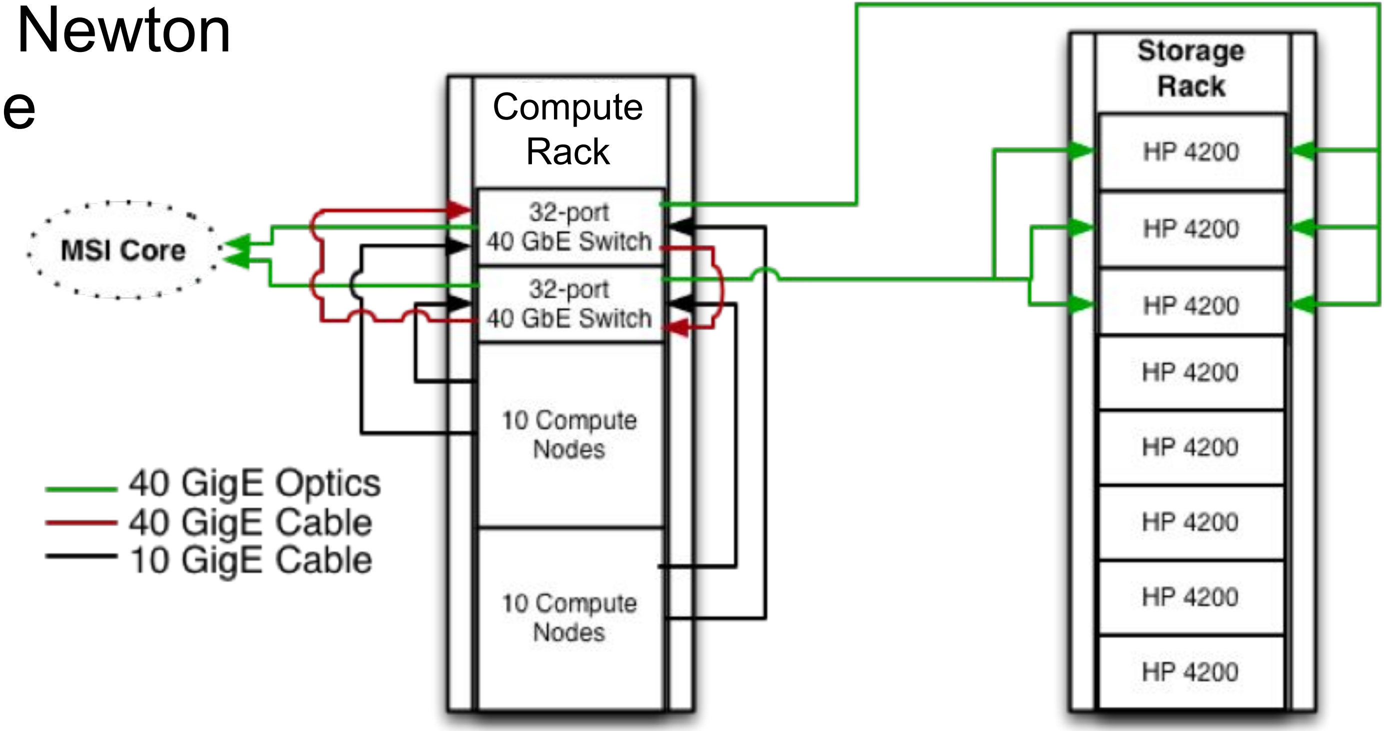

Stratus Features

- Ceph Block Storage
- Ceph S3 Object Storage
- Horizon web interface to manage VMs
- API access

Stratus Features

- Two-factor authentication with University credentials
- MSI-blessed images for common use cases
- Freedom to run user-maintained images
- Subscription model for cost recovery ($165 / month for base subscription)

What Worked?

- Bespoke data compliance and other needs that were not met in the existing HPC cluster ✅

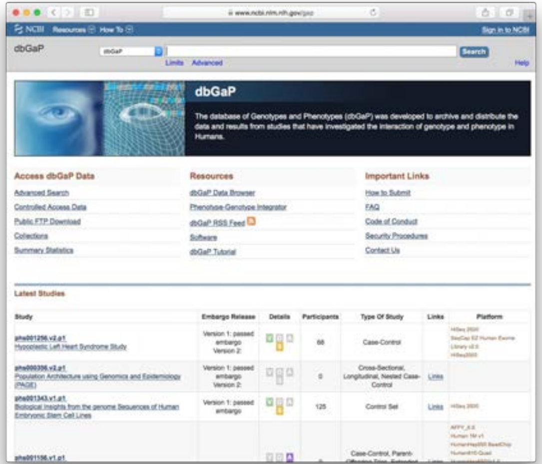

What Worked?

- Bespoke data compliance and other needs that were not met in the existing HPC cluster ✅
- Researcher-controlled persistent services and long-running jobs ✅

Experimental and Persistent User-Controlled Services

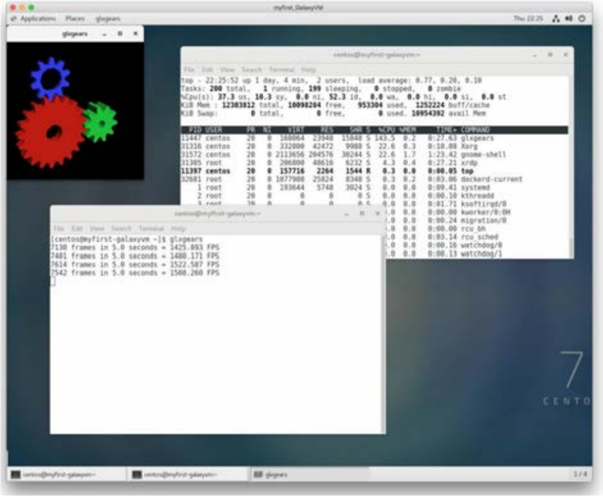

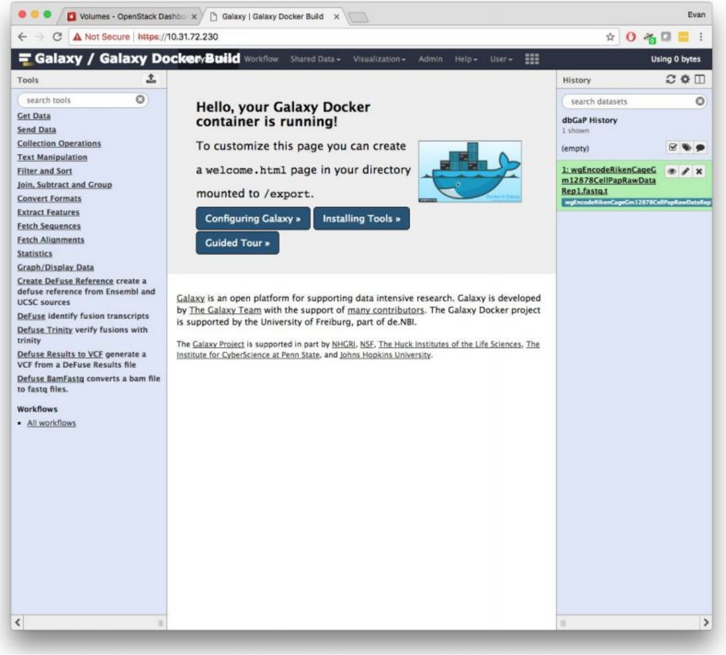

What Worked?

- Bespoke data compliance and other needs that were not met in the existing HPC cluster ✅
- Researcher-controlled persistent services and long-running jobs ✅
- MSI DevOps-Controlled Applications ✅

MSI DevOps-Controlled Applications

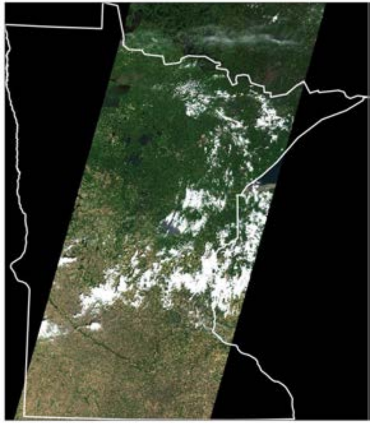

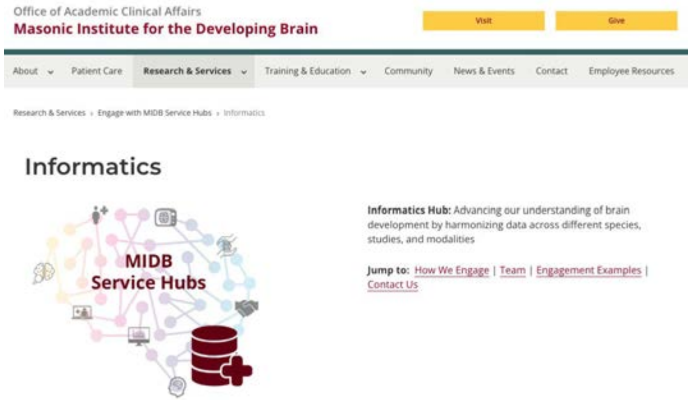

What Worked in 2017?

- Bespoke data compliance and other needs that were not met in the existing HPC cluster ✅
- Researcher-controlled persistent services and long-running jobs ✅
- MSI DevOps-controlled applications ✅

What is Still Working in 2023?

- Bespoke data compliance and other needs that were not met in the existing HPC cluster ✅
- Researcher-controlled persistent services and long-running jobs ✅
- MSI DevOps-controlled applications ✅

What is Still working in 2023?

- Bespoke data compliance and other needs that were not met in the existing HPC cluster
- Researcher-controlled persistent services and long-running jobs ✅
- MSI DevOps-controlled applications ✅

What is Still working in 2023?

- Bespoke data compliance and other needs that were not met in the existing HPC cluster
- Researcher-controlled persistent services and long-running jobs
- MSI DevOps-controlled applications ✅

Future Use Cases

- MSI DevOps-controlled applications
  - Rapidly growing number of applications in neuroscience, GIS, and agroinformatics
  - APIs with calculations that run on GPUs

Future Plans

- Deploy similar-size cluster
- Add GPUs
- Focus on collaborative DevOps projects as the primary use case

Lessons Learned

- Like most of Research Computing, IaaS is a moving target
- Cyberinfrastructure needs require constant reevaluation and communication with researchers.
- Effective use of IaaS generally requires DevOps skills beyond what individual research groups typically maintain.