Lessons Learned in Deploying OpenStack for HPC Users

Graham T. Allan
Edward Munsell
Evan F. Bollig
Minnesota Supercomputing Institute

University of Minnesota
Driven to Discover℠

Stratus: A Private Cloud for HPC Users

Project Goals

- Fill gaps in HPC service offerings
- HPC-like performance
- Flexible environment to handle future needs

OpenStack Cloud

- Multi-tenant
- Self-service VMs and storage

Ceph Storage

- Block Storage for VMs and Volumes
- Additional S3 storage tiers
- Inexpensive to scale

Stratus: Designed for controlled access data

Isolation from MSI Core Services

Two-Factor Authentication

Access Logging

Data-at-rest Encryption

Object storage cache with lifecycle

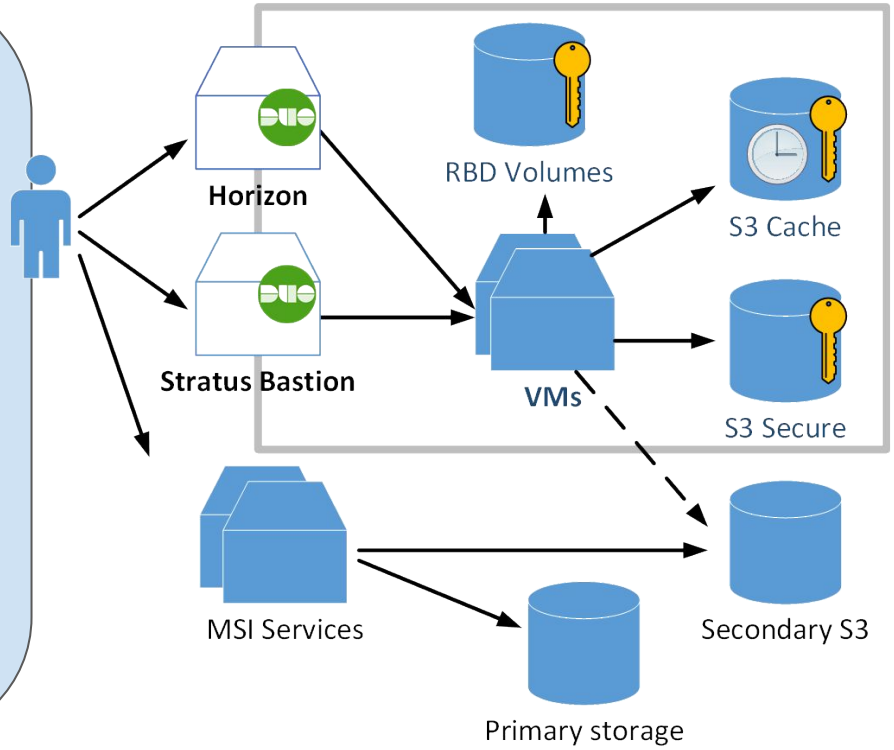

MSI at a Glance

42 Staff in 5 Groups.

4000+ users in 700 research groups.

Major focus on batch job scheduling in a fully-managed environment.

Most workflows run on two HPC clusters.

Traditional HPC — "Big Data"

Physical Sciences — Life Sciences

Mesabi cluster (2015)

Haswell-based, 18000 cores, memory sizes 64GB, 256GB & 1TB

Some specialized node subsets: K40 GPUs, SSD storage

800 TFLOPs, 80TB total memory

Still in top-20 of US University-owned HPC clusters

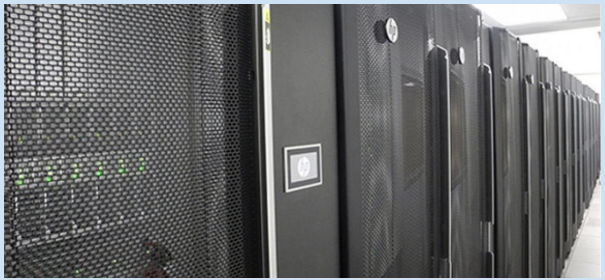

MSI at a Glance

Allocated CPU hours vs Discipline

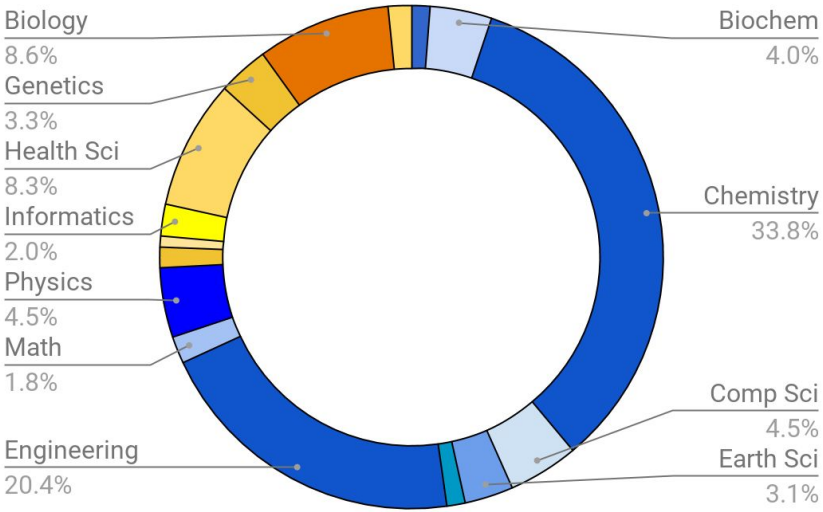

Allocated storage vs Discipline

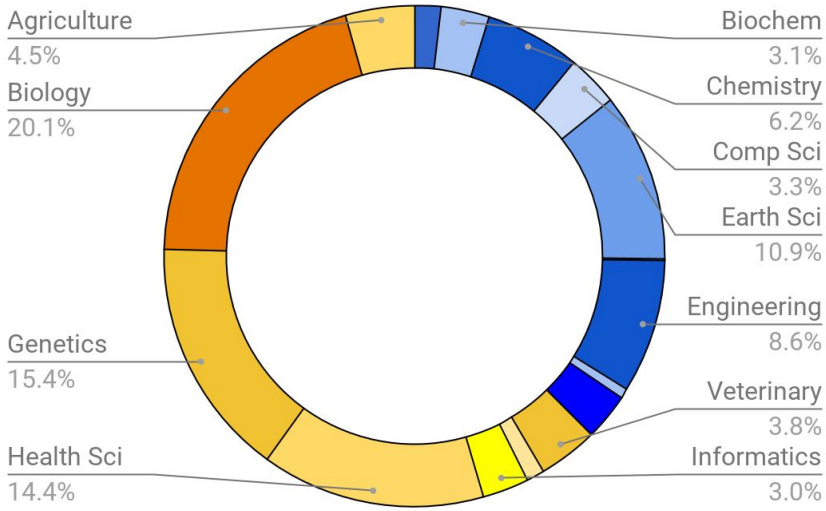

Life sciences consume only 25% of cpu time but 65% of storage resources

Physical sciences consume 75% of cpu time but only 35% of storage.

Stratus: Why did we build it?

#1 Environment for controlled-access data

#2 On-demand computational resources

#3 Demand for self-service computing

#4 Satisfy need for long-running jobs

Intended to complement MSI's HPC clusters, rather than compete with them...

Controlled-access data

dbGaP: NIH Database of Genotypes and Phenotypes

40+ research groups at UMN.

Data is classified into "open" and "controlled" access.

"Controlled access" governed by Genomic Data Sharing policy

Requires two-factor authentication, encryption of data-at-rest, access logging, disabled backups… etc…

Standard HPC cluster gives limited control over any of these.

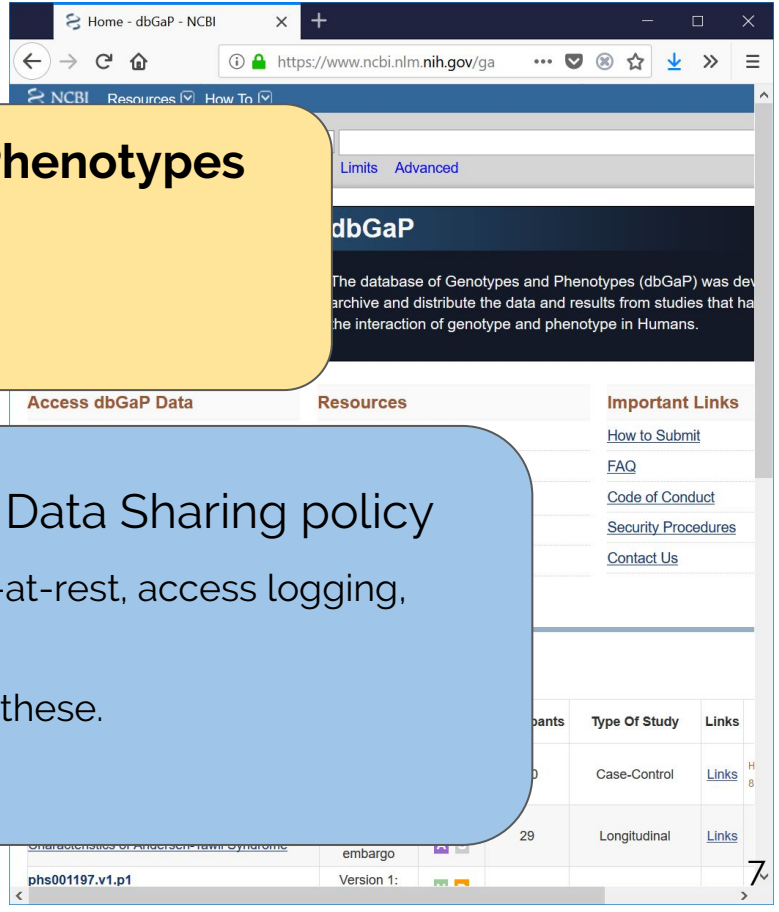

Controlled Access Data: Explosion in Size

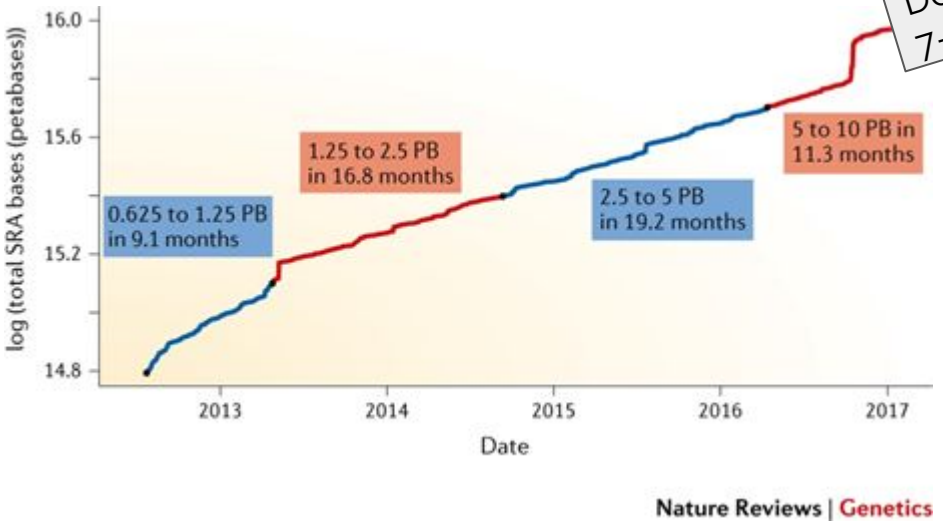

Increase in storage of genomic sequencing data (estimated **8 Petabytes in 2018**)

Reprinted by permission from: Macmillan/Springer, Nature Reviews Genetics, Cloud computing for genomic data analysis and collaboration, Ben Langmead & Abhinav Nellore, 2018

Doubling every 7-12 months

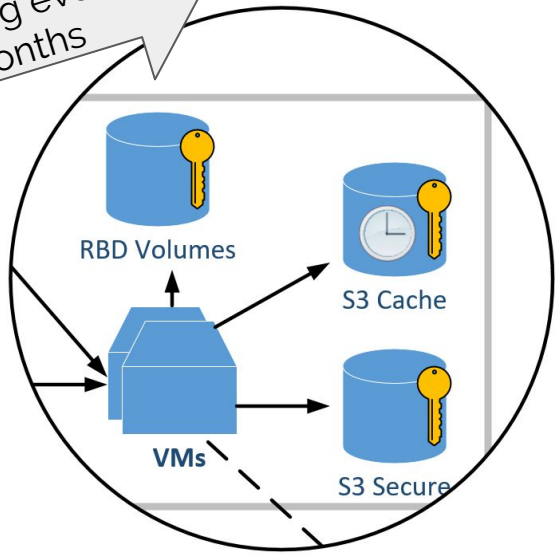

Cache model for stratus object store based on this large & increasing data size

MSI not a NIH Trusted Partner: no persistent copy of data.

Expanding the scope of Research Computing

Should MSI be the home of such a project, vs some other organization?

Discussion of MSI's evolving role in supporting research computing.

Existing culture based on providing fully-managed HPC services.

Fear that self managed VMs could undermine infrastructure security.

Weekly “Best Practices for Security” meeting (Therapy sessions).

Working with controlled-access data was previously discouraged.

Focus on dbGaP-specific data controls and avoid scope creep.

Timeline

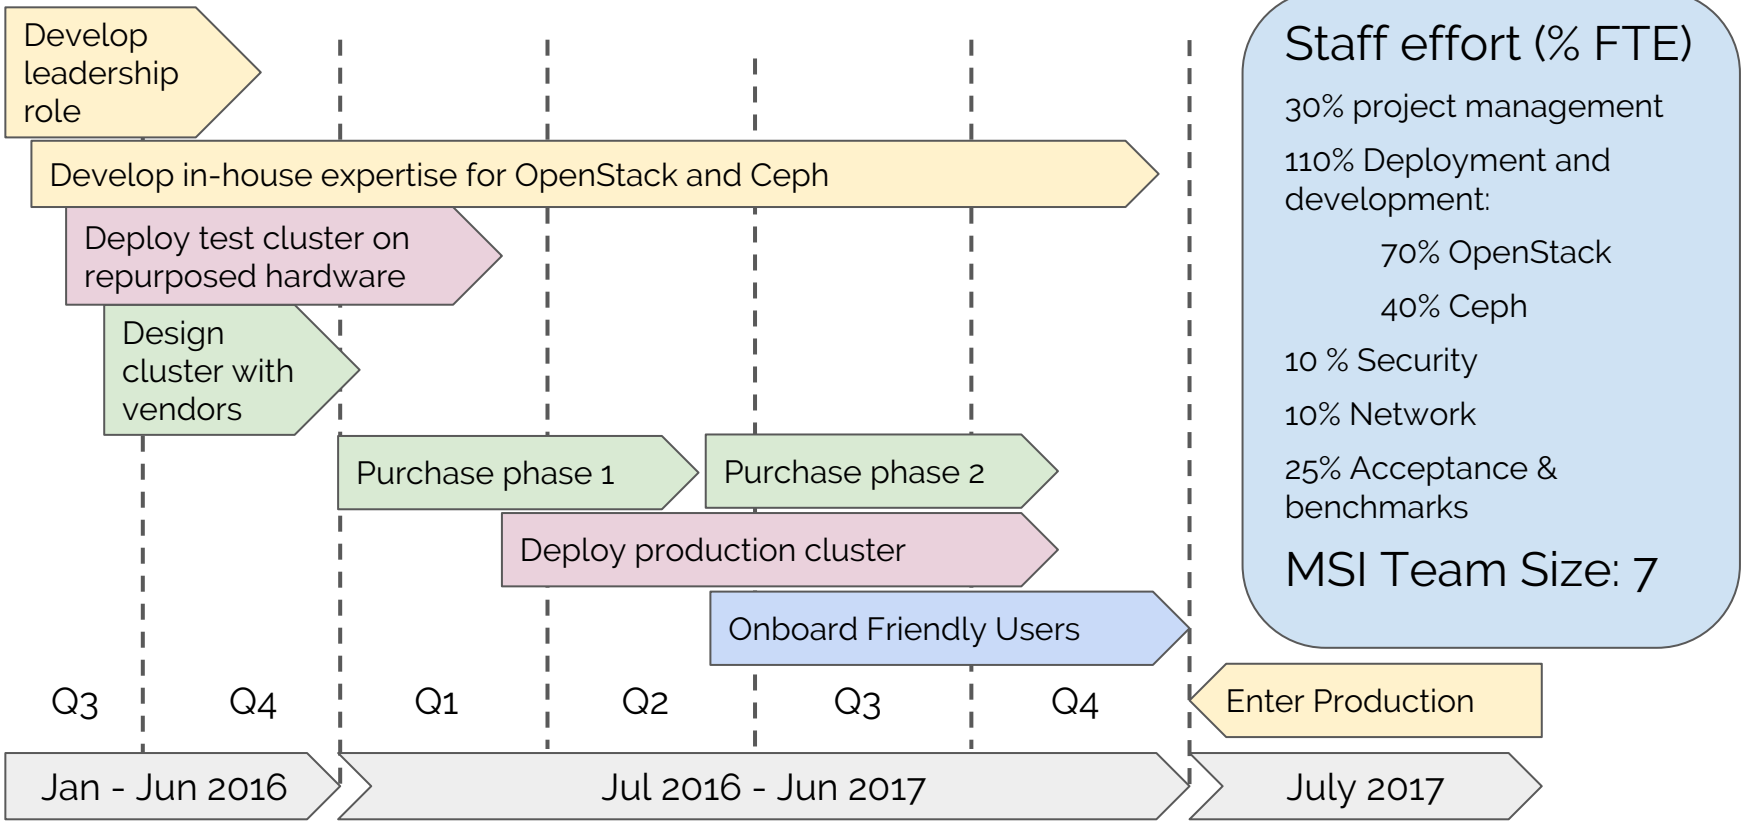

Staff effort (% FTE)

30% project management

110% Deployment and development:

- 70% OpenStack
- 40% Ceph

10 % Security

10% Network

25% Acceptance & benchmarks

MSI Team Size: 7

Development Cluster

OpenStack and Ceph components

Develop hardware requirements

Gain experience configuring and using OpenStack

Test deployment with puppet-openstack

Test ability to get HPC-like performance

Cloud vs Vendor vs Custom Solution

Cloud solutions

Performance and scalability - relatively high cost

Discomfort with off-premises data

Vendor solutions

Limited customization

Targeted to enterprise workloads, not HPC performance

Not cost effective at needed scale

Custom OpenStack deployment

Develop in-house expertise

Customise for security and performance

Resulting Design

20x Mesabi-style compute nodes

- HPE Proliant XL230a
- Dual E5-2680v4. 256GB RAM
- Dual 10GbE network
- No local storage (OS only)

8x HPE Apollo 4200 storage nodes

- 24x 8TB HDD per node
- 2x 800GB PCIe NVMe
- 6x 960GB SSD
- Dual 40GbE network

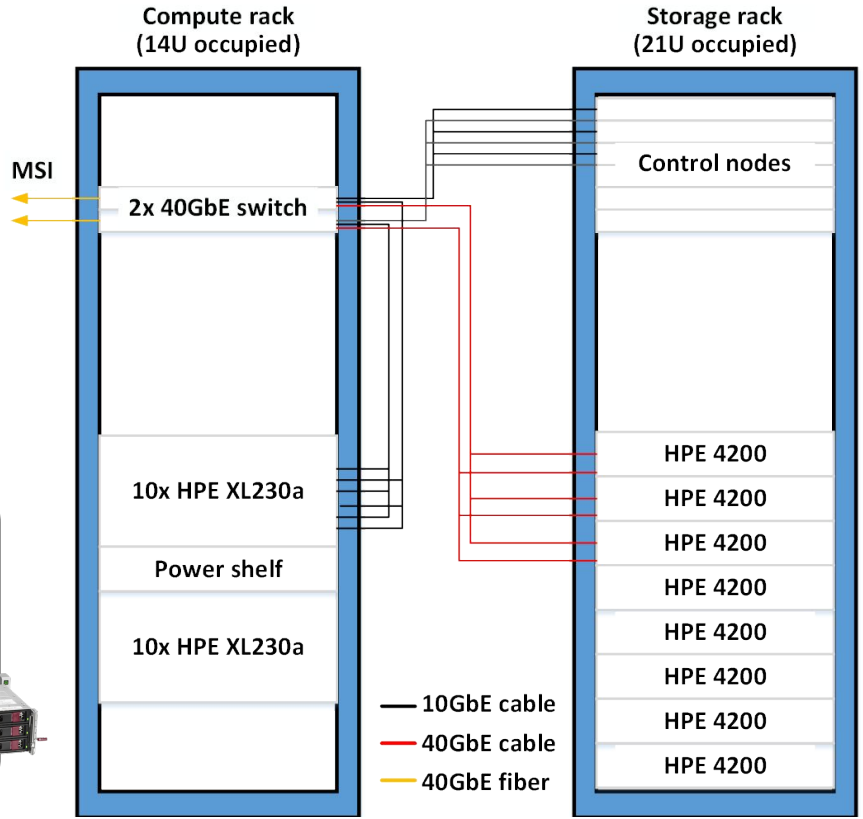

Resulting Design

10x support servers

Repurposed existing hardware...

Minor upgrades of CPU, memory, network, as work-study projects for family members.

Controllers for OpenStack services.

Ceph mons and object gateways.

Admin node, Monitoring (grafana).

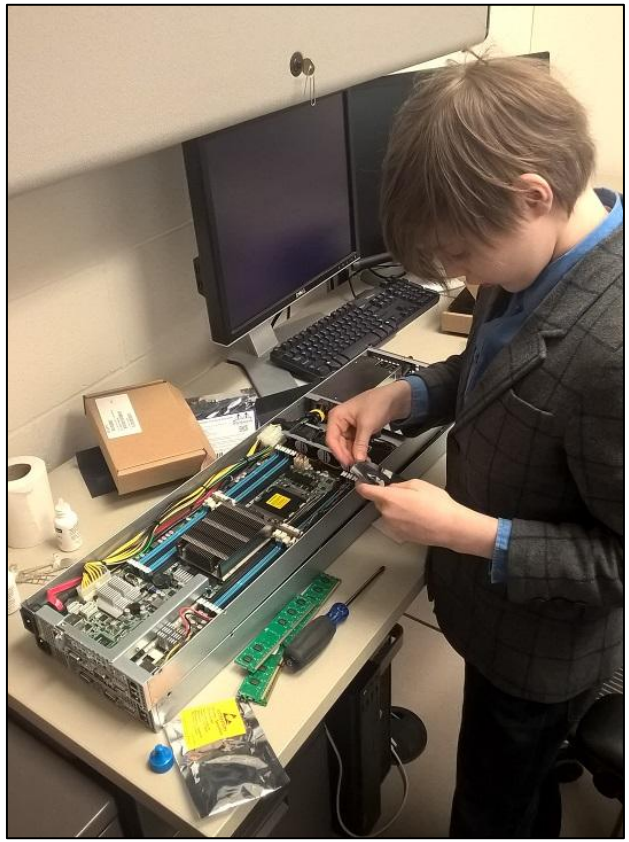

Stratus: OpenStack architecture

Minimal set of OpenStack services

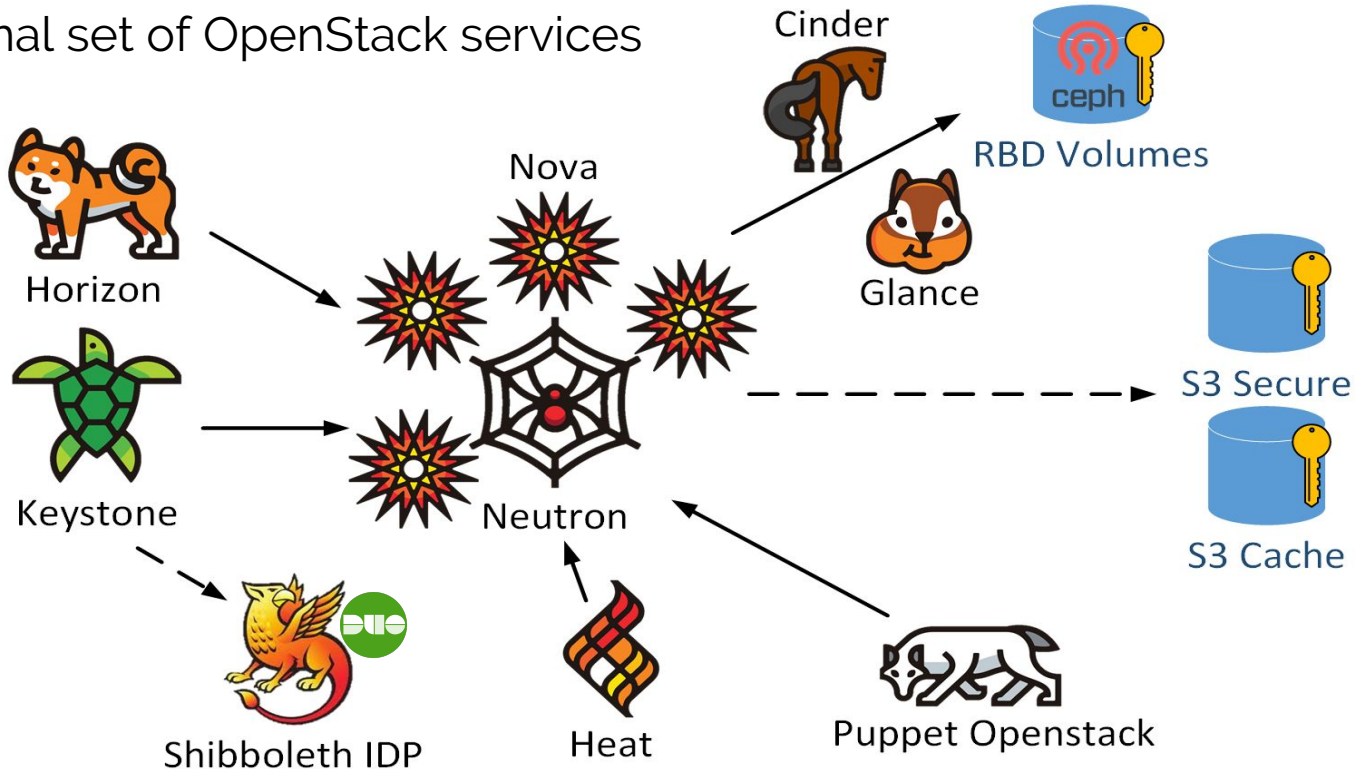

Stratus: Storage architecture

Eight HPE Apollo 4200 storage nodes

HDD OSDs with 12:1 NVMe journals: 1.5PB raw

- 200GB RBD block storage, 3-way replicated
- 500GB s3 object storage, 4:2 erasure coded

SSD OSDs: 45TB raw

- object store indexes, optional high speed block

Configuration testing using CBT

- Bluestore vs Filestore
- NVMe journal partition alignment
- Filestore split/merge thresholds
- Recovery times on OSD or NVMe failure
- LUKS disk encryption via ceph-disk: <1% impact

Bluestore vs filestore read performance

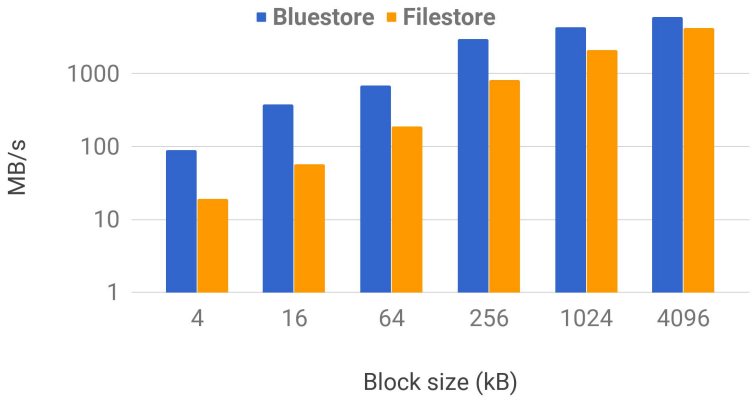

HPC-like performance

HPL Benchmark

Popular measurement of HPC hardware floating point performance.

Stratus VM results

95% of bare-metal performance

CPU-pinning and NUMA awareness disabled

Hyperthreading, 2x CPU oversubscription

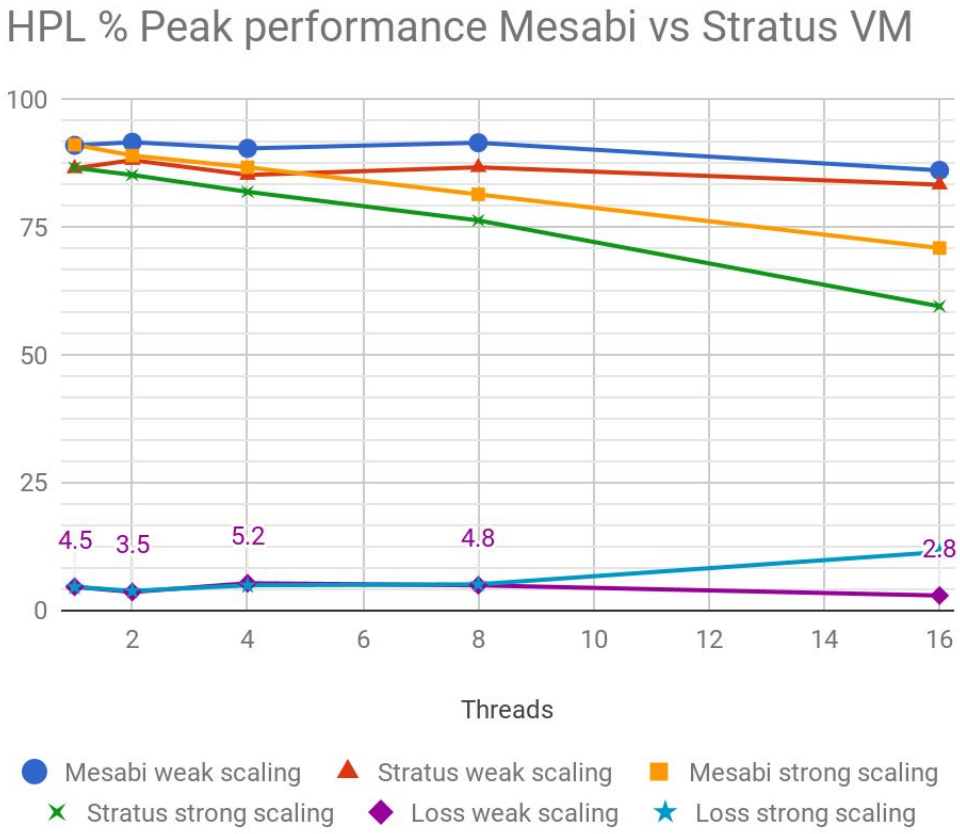

HPC-like storage

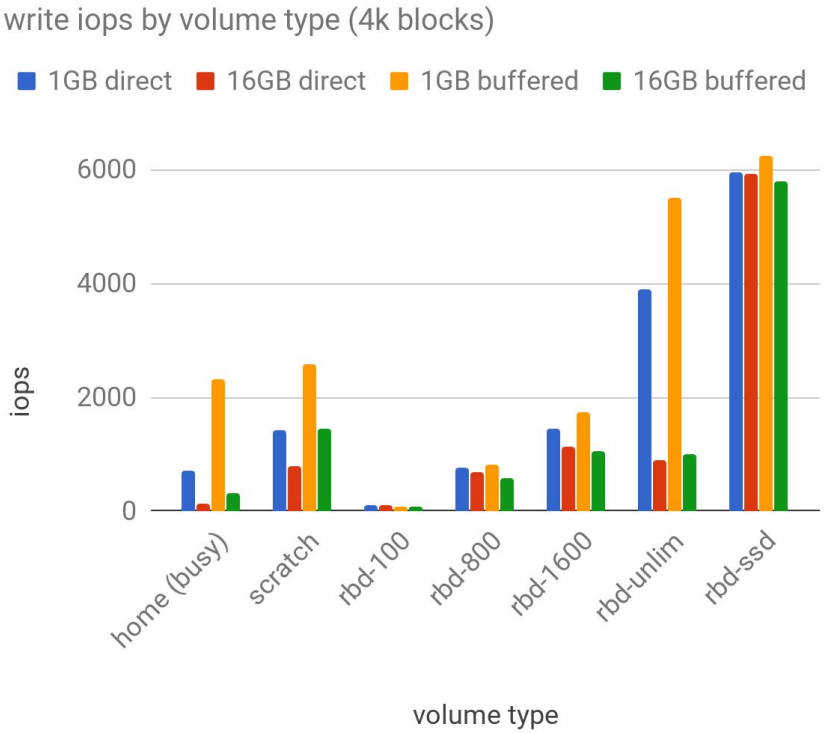

> "**We claim that file system benchmarking is actually a disaster area** - full of incomplete and misleading results that make it virtually impossible to understand what system or approach to use in any particular scenario."
>
> File System Benchmarking: It *IS* Rocket Science, Usenix HOTOS 11, Vasily Tarasov, Saumitra Bhanage, Erez Zadok, Margo Seltzer

Select benchmark: *FIO - mixed read/write random iops*

Characterise storage performance for Mesabi single node

Characterise performance on Stratus for single and multiple VMs.

Dial-in default volume QoS limits to provide close match to Mesabi, balanced against scalability.

User Experience Preview

Staff performing benchmarks & tests expected a managed HPC environment.

Non-sysadmins managing infrastructure for the first time

- No scheduler or batch system
- No pre-installed software tools
- No home directory
- Preview of pain points for regular users

Bringing in our first Users

Users excited by freedom and flexibility expected from a self-service environment

... but are shocked to discover what is missing.

Introductory tutorial

- Introduce security measures and shared responsibilities
- Introduction to OpenStack, how to provision VMs and storage
- Crash course in basic systems administration

Recurring questions

- Where is my data and software?
- How do I submit my jobs?
- Who do I ask to install software?

Pre-configured Images

dbGaP "Blessed"

CentOS 7 base preloaded with GDC data transfer tools, s3cmd and minio client, Docker, R, and a growing catalog of analysis tools

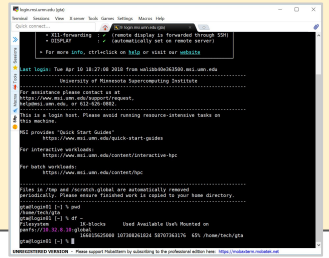

Blessed + Remote Desktop

RDP configured to meet security requirements: SSL, disable remote media mounts.

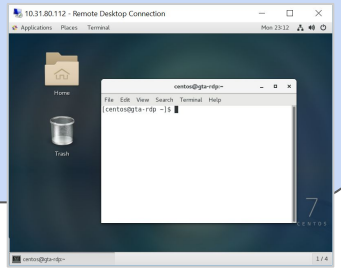

Blessed + Galaxy

Galaxy is a web-based tool used to create workflows for biomedical research

Shared responsibility security model

Genomic Data Sharing policy as a good starting point

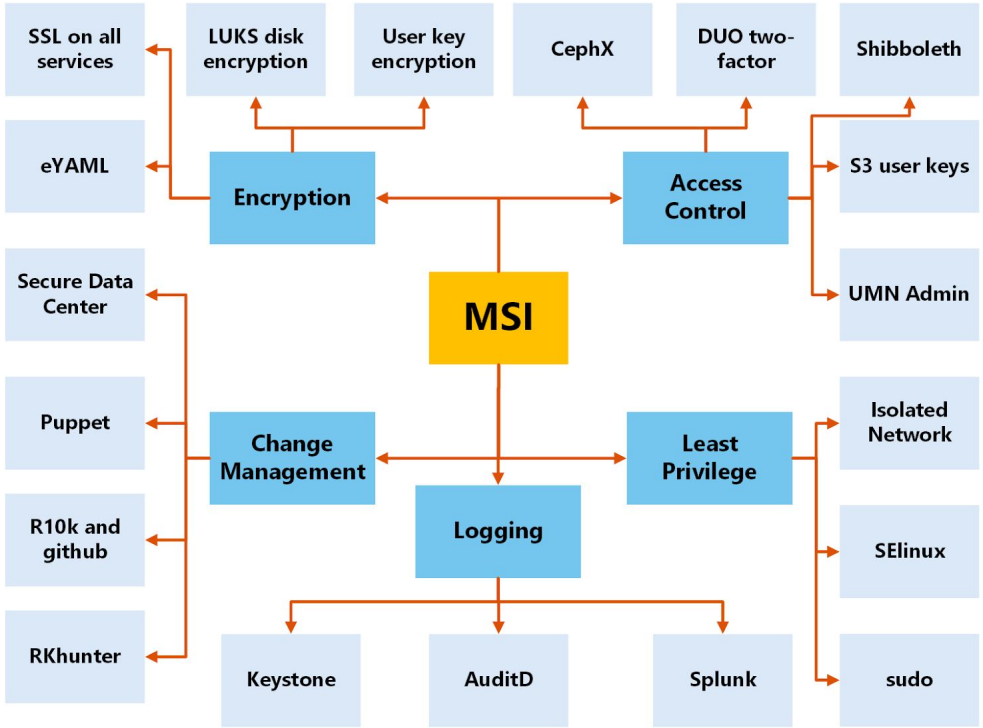

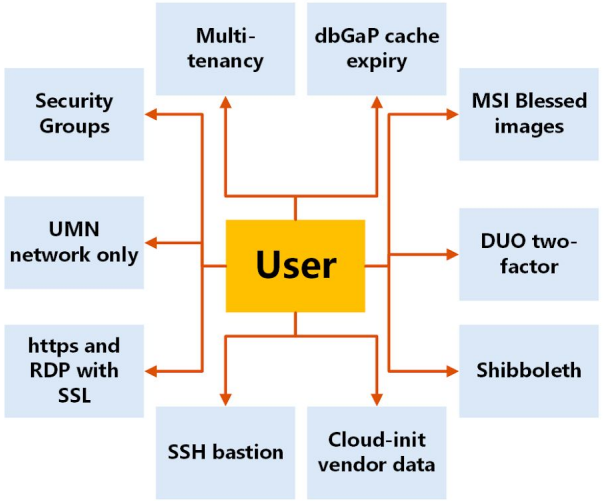

Left shows controls on MSI infrastructure
Right shows controls on user environment

Security Example: Network isolation

Campus network traffic only

https and rdp ports only

SSL-encryption required.

Tenants cannot connect to other tenants

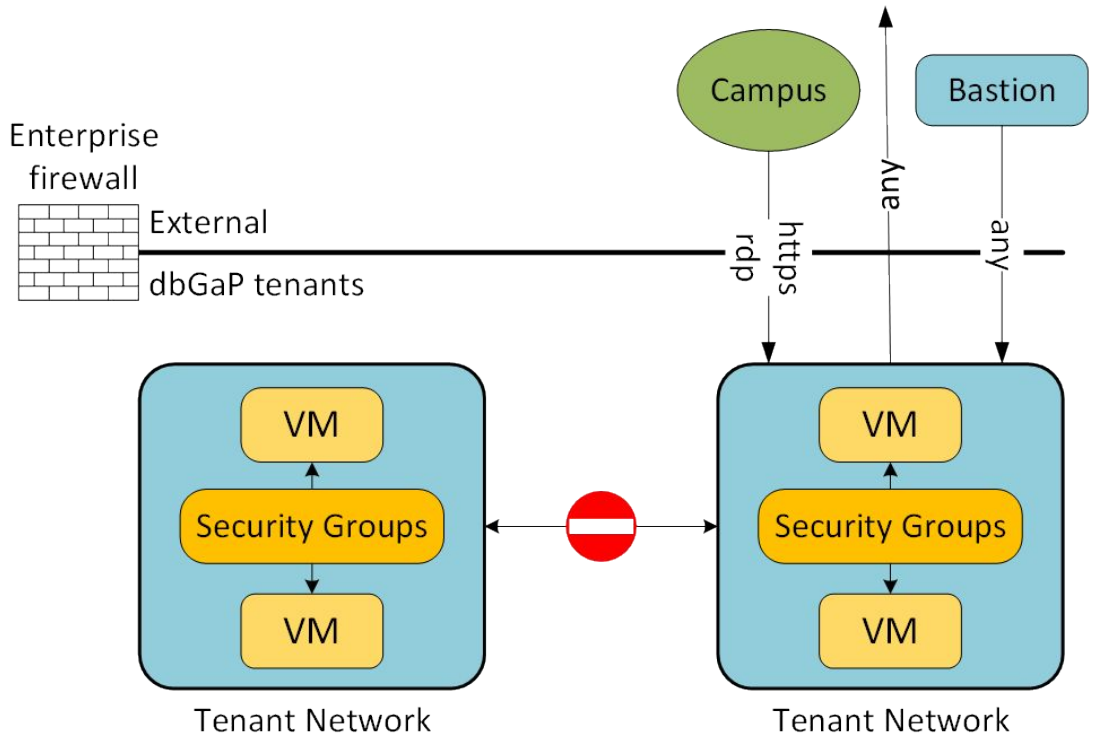

Cost Recovery

Stratus introduced as a subscription service

- Discourage superficial users
- Zero profit
- Build in staff FTE costs for support
- Base subscription with a la carte add-ons.
- Target 100% of hardware cost recovery at 85% utilization

Cost comparison showed AWS to have significantly higher costs (11x) for equivalent subscription.

Annual base subscription: $626.06 (internal UMN)

- 16 vCPUs, 32GB memory
- 2TB block storage
- access to 400TB S3 cache

Add-ons:

vCPU + 2GB memory: $20.13

Block storage: $151.95/TB

Object storage: $70.35/TB

Lessons from Production #1

Users are willing to pay for convenience

On the first day Stratus entered production, our very first group requested an extra 20TB of block storage (10% of total capacity)

Users are accustomed to POSIX block storage and willing to pay for it.

We increased efforts to promote using the free 400TB s3 cache in workflows.

But object storage is still alien to many users.

Lessons from Production #2

Layering of additional support services

Initially started with a ticket system for basic triage

Some users hit the ground running

Some needed more help...

Additional (paid) consulting options:

From Operations

- setup or tuning of virtual infrastructure

From Research Informatics group

- help develop workflows
- perform entire analysis

Lessons from Production #3

Heavier demands came sooner than expected

New research group with much larger resource needs.

Working on whole-genome (TOPMed) data - 100x larger than exome.

Used to running on 1TB HPC cluster nodes.

Need for multiple VMs with 50 cores, 100-200GB memory.

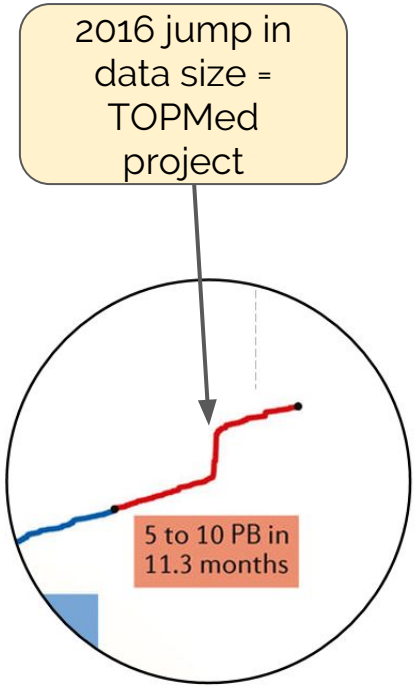

Lessons from Production #4

Users universally asked for a more flexible subscription model

Changed subscription from annual to quarterly.

	Annual	Quarterly
Base Subscription 16 vCPU, 32GB RAM, 2TB block	$626.06	$156.52
Additional vCPU with 2GB RAM	$20.13	$5.04
Block storage per TB	$151.95	$37.99
Secure object storage per TB	$70.35	$17.59

Lessons from Production #5

Expand access beyond dbGaP users

Added a new "general" provider to meet additional use cases (February 2018)

Open network access from campus

No access to the secure object stores

#2 On-demand resources

#3 Self-service computing

#4 Long-running jobs

Conclusions

Did we make a good decision?

For MSI...?

Issue of securely handling controlled access data had to be addressed

Stratus gives a solid starting point to expand to other sets of requirements (eg FISMA, FedRAMP)

For our users...?

Stratus does provide performance, security and flexibility for them to build a successful research environment

But, their lives have become more complicated. Some diversity in ease of adaptation.

Conclusions

Would we build it the same way again?

Custom environment provided flexibility and scale which vendor solutions couldn't match

Strength of OpenStack community in solving problems

What would we change?

s3 object cache is an elegant technical solution but is underutilized - roadblock for user workflows.

Future Work

Manage user encryption keys with Barbican

- Help users meet dbGaP requirement for encryption using user keys
- Easier user encryption of S3 data
  - We currently recommend using minio client with SSE-C
  - SSE-KMS with Barbican probably more transparent
- User encryption of cinder volumes

Heterogeneous nodes

- Requirements for large memory systems (1TB)
- Virtual GPUs for machine learning users

Future Work

Storage as a Service

- Desire for shared POSIX storage between multiple VMs
- Multi-attach RBD volumes (read-only)
- Manila NFS volumes

HPC as a Service

- Some users struggle with lack of job control

Thank You

Any Questions?

HPC-like performance

HPL Benchmark

HPL weak scaling

Stratus bare-metal vs Stratus VM (28 vCPU)