From Bare Metal to Virtual: Lessons Learned when a Supercomputing Institute Deploys its First Cloud

Evan F. Bollig (bollig@umn.edu)
James C. Wilgenbusch
Minnesota Supercomputing Institute

http://hdl.handle.net/2022/21721

University of Minnesota
Driven to Discover℠

Background (What's in a Name?)

1980's -- "Supercomputing"

1990's -- "High Performance Computing (HPC)"

2000's -- "Research Computing"

Cyberinfrastructure used to support research

Research supported by cyberinfrastructure

Since 1984, MSI has supported computationally intensive research

- Academic unit under the Office of the Vice President for Research (Research Computing umbrella)
- 42 staff (5 functional groups) for 700 groups, 4000+ users
- Supports public and private entities throughout Minnesota

Background (MSI Core Services)

Homogenous environment simplifies experience and satisfies most data-use agreements

- Most workflows generalize to large HPC clusters (Mesabi and Itasca)
- Tiered storage with global namespace
- Central OIT ID management and authentication

Edge cases handled as one-offs

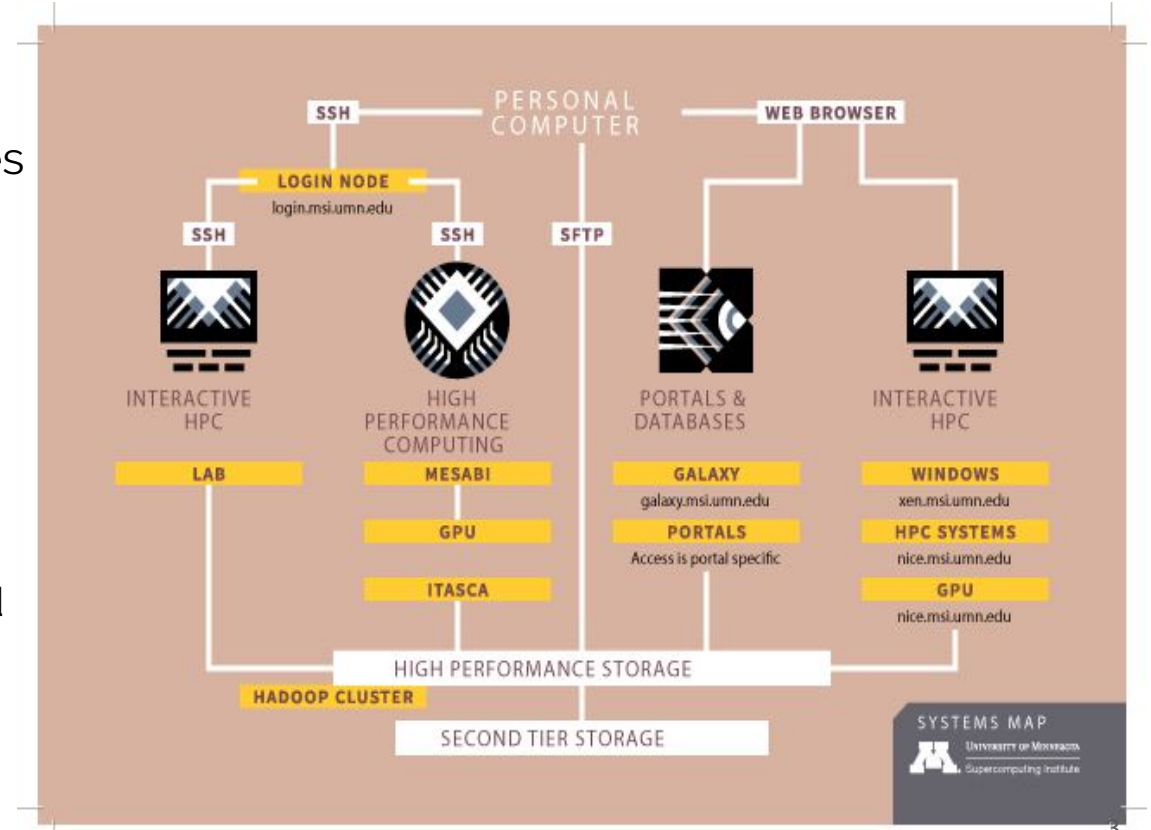

Requirements and Planning (Unmet Needs)

Edge-Cases become cumbersome

- On-demand resources for data intensive research
  - 10s - 100s TBs of data
  - Non-traditional HPC software and workflows
- Long running jobs
  - Monthly maintenance day limits jobs to 29 days
- Container-based computing
  - Docker and sudo pose security risks
- NIH Controlled-Access Data (dbGaP)
  - Limited control over authentication, isolation, and logging
  - Growing number of researchers (40+)
  - One-off model for backups, access control, etc. does not scale
  - Size and cost of storage is high with unique copies required

Is MSI the right provider for these?

A: Yes. With a new resource.

Requirements and Planning (Proposed Solution)

- OpenStack Cloud
  - Self-service VMs and Volumes
  - Containers within VMs
  - Live Migration for long jobs
- Ceph Storage
  - Inexpensive to scale
  - Block Storage for VMs and Volumes
  - Object Storage for secure S3 Cache and Persistent Storage
- Designed for NIH dbGaP data
  - Two-factor Authentication
  - Encryption and Access Logging
  - Isolation from Core Services
  - 60 lifecycle on S3Cache

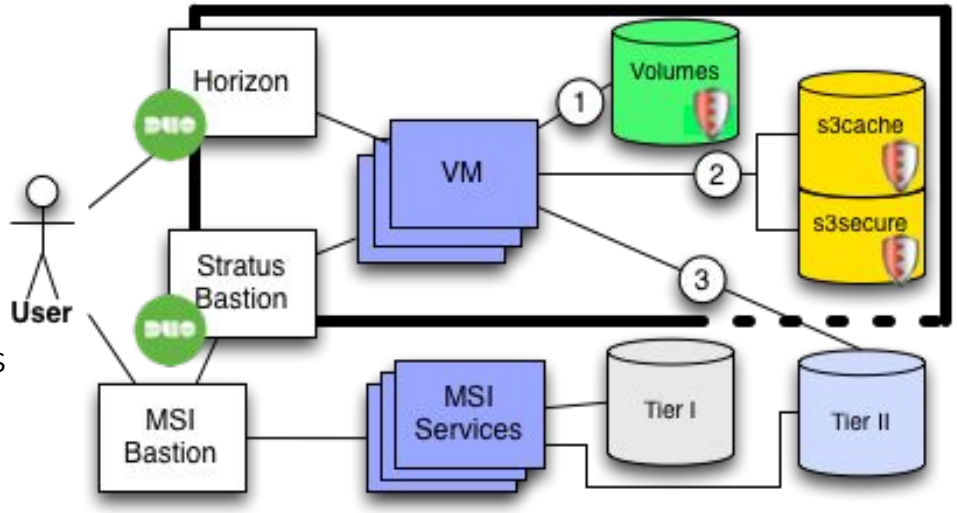

Stratus Compute Cloud

An Ambitious Timeline

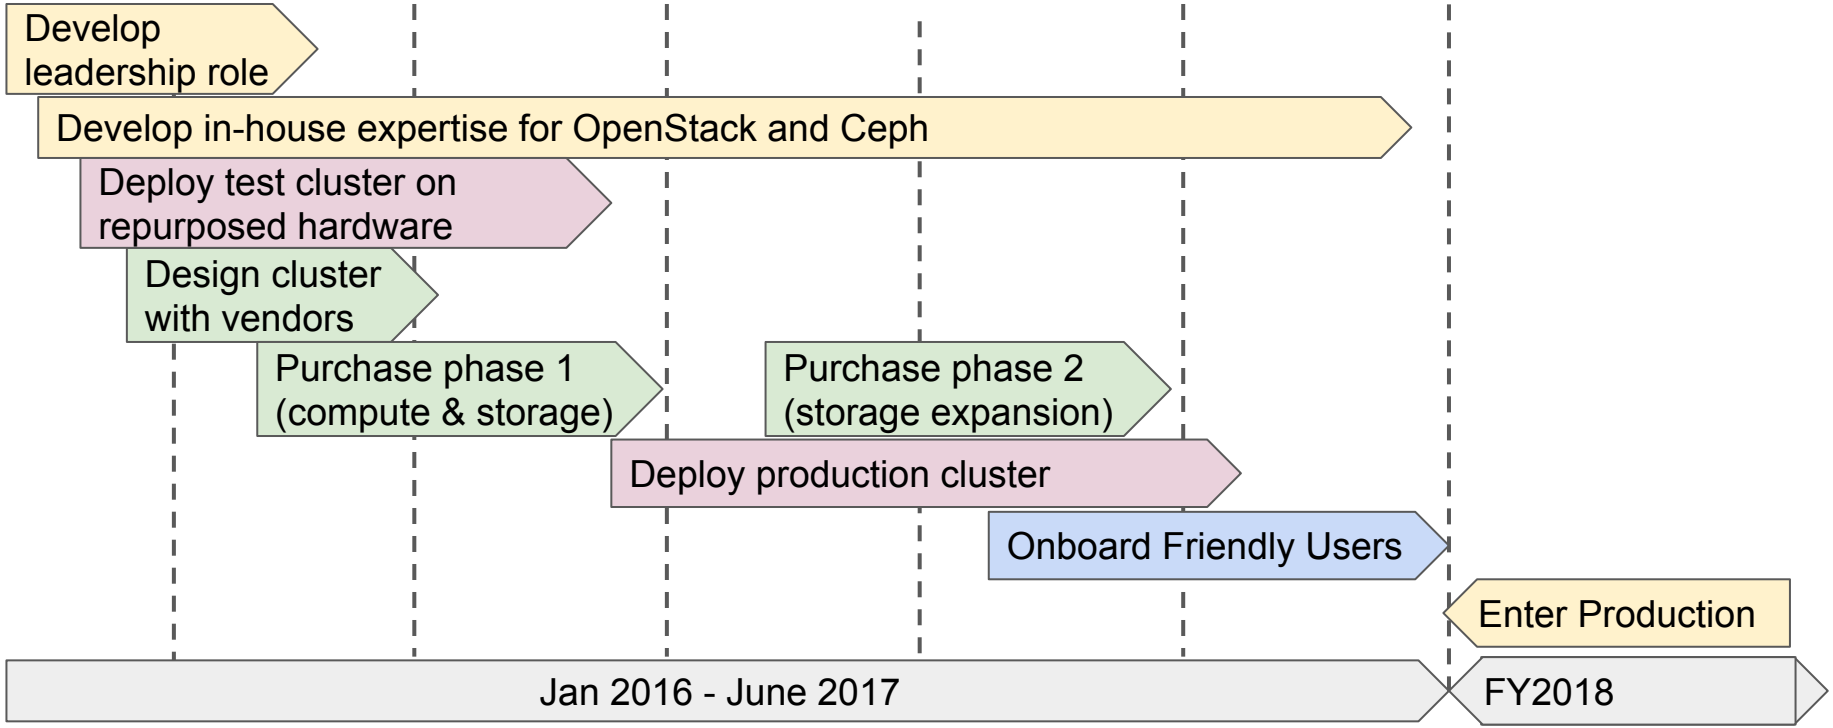

Training and Staff Culture (Responsibilities)

New Leadership Roles

- MSI restructured as a Matrix
  - 5 functional verticals (groups)
  - A project spans multiple verticals
- Project Manager
  - Backed by functional leads
  - See project through operational hand-off
- Project Staff
  - Reports to functional lead first
  - Time allocated to project

Staff Effort (% FTE):

- OpenStack deployment, development -- 70%
- Ceph deployment -- 40%
- Acceptance tests and benchmarks -- 25%
- Security -- 10%
- Network -- 10%

MSI Team size: 7

Training and Staff Culture (Ownership)

Lesson Learned: For staff, finger pointing is easier than ownership.

Solution: Take ownership of the project, and compliance. Lead by example.

- Defend logic behind MSI--not another dept or org--as the choice to build this research-centric service.
- “Therapy session” to reassure staff that technical design and documentation will not cost them a job, $$, or prison time.
- Weekly “Best Practices for Security” meeting to demystify standards, and open the dialog about implementation.

Training and Staff Culture (Staff-Development)

Lesson Learned: Staff do not always share the same appreciation for adding new services (a.k.a. “responsibilities”). Expect pushback.

Solution: Emphasize Professional Development Opportunities

- Research Experience and Co-authorship
  - Four first-time authors on our first submission (https://doi.org/10.1145/3093338.3104185)
- OpenStack and Ceph are hot new skillsets
  - OpenStack Cloud Engineer salaries are **36% higher** than industry average
- Cross-training for storage, network, automation, etc.

System Configuration, Acquisition, and Installation

Lesson Learned: Vendor solutions were rigid (i.e., no customization), did not meet security needs and would not be cost effective at scale of dbGaP.

Solution: Custom OpenStack deployment on compatible hardware.

- Develop in-house experience
- Cut costs
- Satisfy Requirements

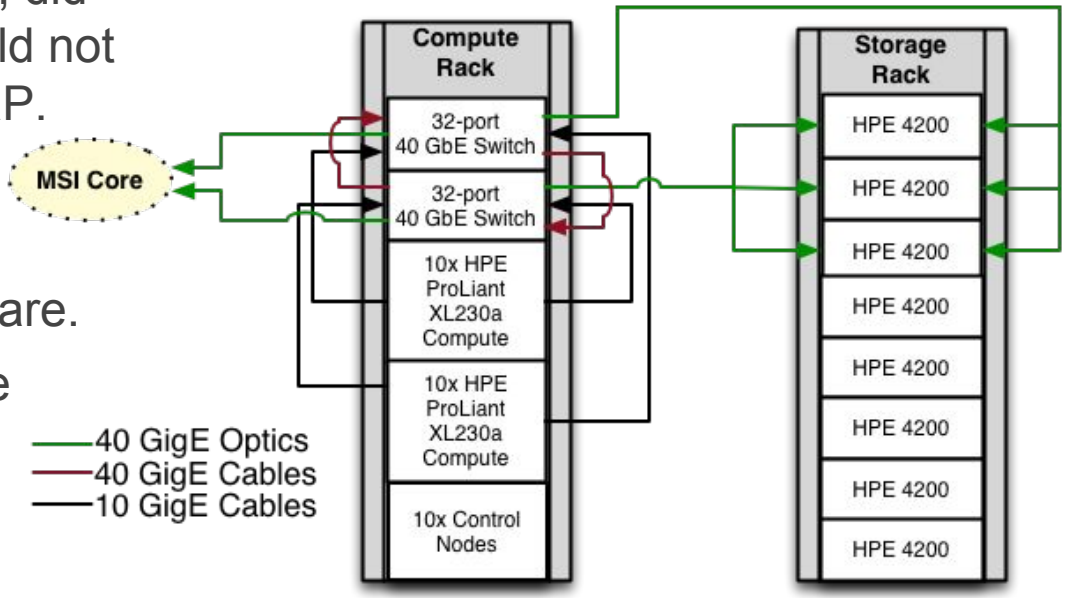

Security Planning (MSI-first Mentality)

Lesson Learned: the NIH GDS Policy is fairly lax, but a good launching point.

Solution: Use the NIH dbGaP Best-Practices Guide as a checklist. Expand to more stringent data policies later.

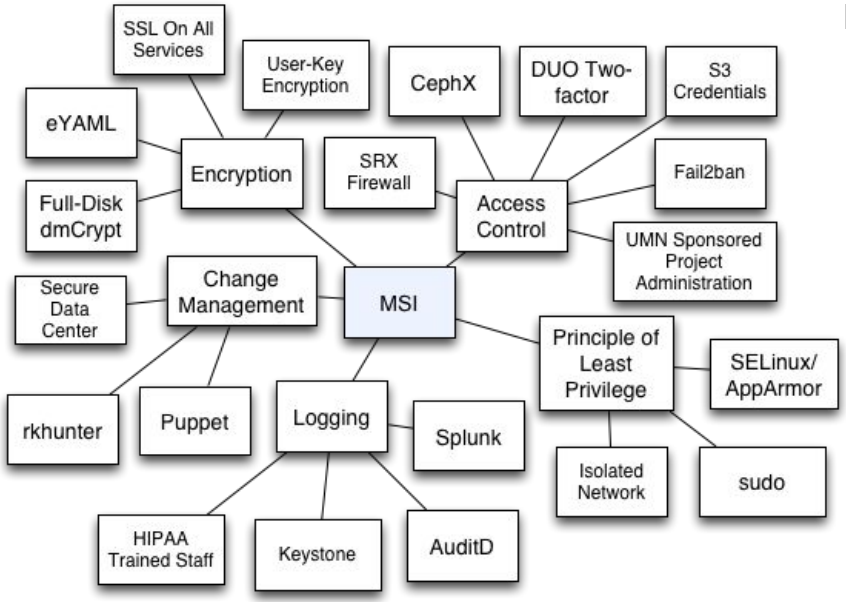

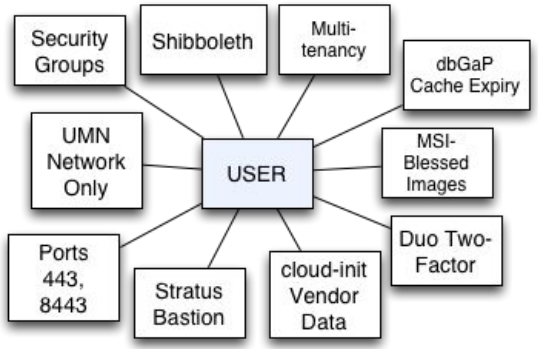

Security Planning (Example: Isolating Users)

Stratus only allows campus network traffic on ports and 443, and 8443 with SSL-encryption required.

Tenants cannot connect to other tenants

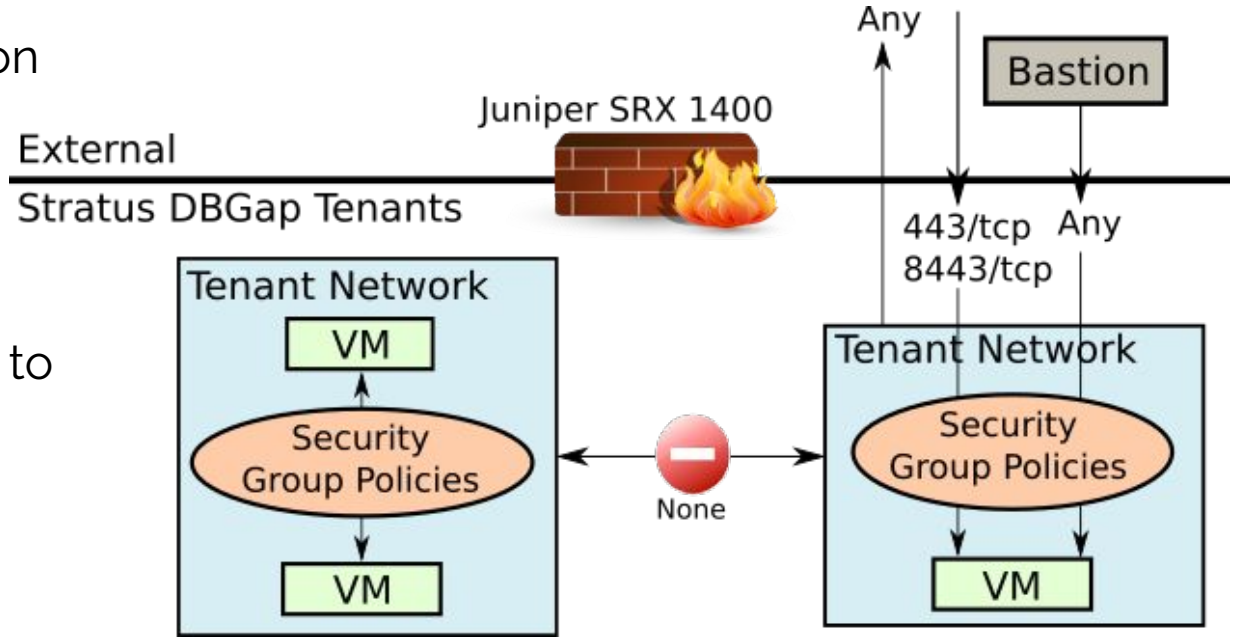

Security Planning (Another Example: Base Images)

We provide a number of base images:

- Vanilla
- dbGaP Blessed
- dbGaP Blessed with Remote Desktop

Lesson Learned: Staff and users will use pre-configured options like Galaxy, and Remote Desktop.

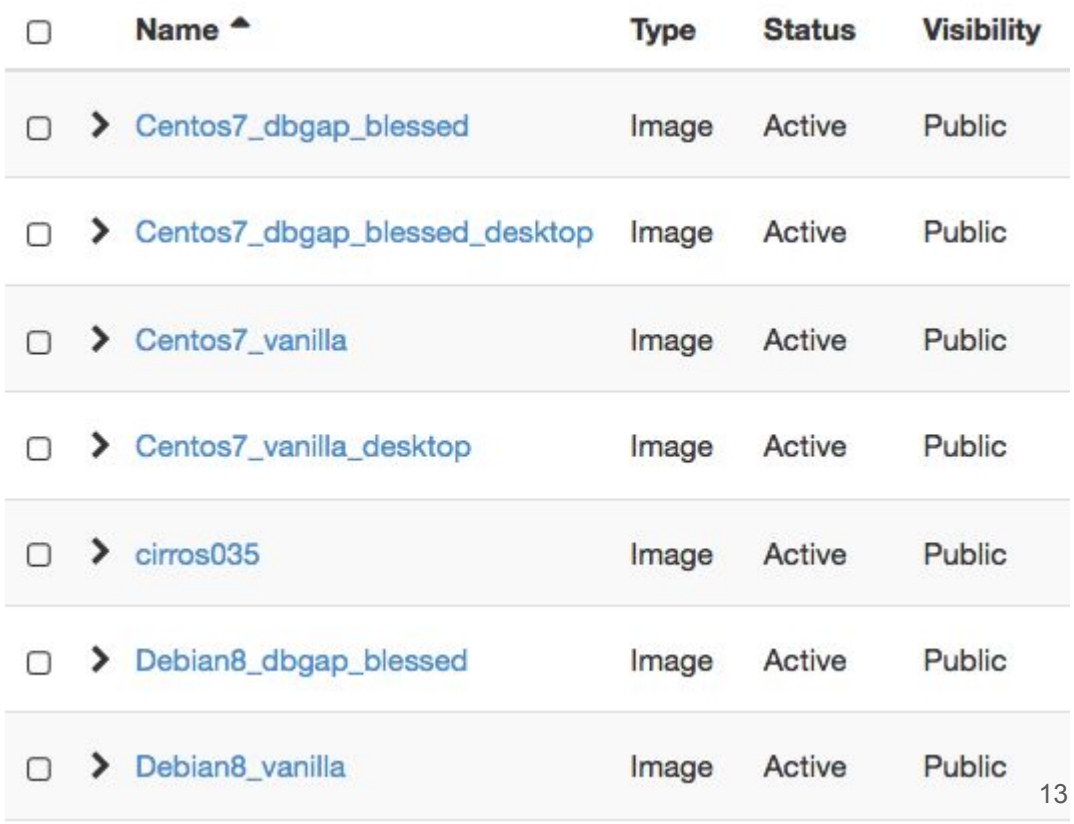

Name	Type	Status	Visibility
Centos7_dbgap_blessed	Image	Active	Public
Centos7_dbgap_blessed_desktop	Image	Active	Public
Centos7_vanilla	Image	Active	Public
Centos7_vanilla_desktop	Image	Active	Public
cirros035	Image	Active	Public
Debian8_dbgap_blessed	Image	Active	Public
Debian8_vanilla	Image	Active	Public

Systems Installation and Testing

Lesson Learned: Network, Compute, Storage, it's all there right?

- A single friendly dbGaP user pulled 120TB of data
  - Storage expansion (+1PB raw) was necessary even before first release

Lesson Learned: Staff performing benchmarks and tests expected a managed HPC environment and pre-installed software.

- Train everyone to self-manage infrastructure for the first time
  - Expect a similar pain-point with regular users

Lesson Learned: Benchmarks revealed 5% efficiency loss between bare-metal and virtualization.

Cost Recovery (Zero-Profit Model)

- Include all hardware purchased
  - Target 100% recovery at 85% utilization
- Build in staff FTE costs for support (administration, ticket triage, training, etc.)
- Structure as annual subscription with a la carte extensions
  - 16 vCPUs, 2TB block storage, 32GB memory, and access to 500TB S3Cache

Service Name	Unit	Cost/Year
Stratus base subscription	Pkg	$626.06
Additional CPU Cores	vCPU	$20.13
Additional Block Storage	TB	$151.95
Persistent Secure Object Storage	TB	$70.35

Lesson Learned: private clouds are significantly cheaper than public cloud

Onboarding Users (Identifying Users)

Lesson Learned: dbGaP users exist, but no one has a complete list of them.

1. University's Sponsored Project Administration (SPA) must approve PI for dbGaP project
2. PI chooses individuals who are granted access to data on the project
3. MSI depends on both SPA and PI to share authorizations
    a. SPA notifies MSI of new PI
    b. Email to PI advertises Stratus as option to store and analyze dbGaP data
    c. PI initiates subscription request and reports list of users
    d. MSI reviews project status annually to renew the subscription (obeying expiration dates is PI responsibility)

Onboarding Users (Training)

Lesson Learned: users don't really understand what they're asking for when they demand self-service.

First questions during onboarding:

1. "How do I run jobs?"
2. "Where is my data and software?"
3. "Where do I send requests for software installs, or system administration?"

Solution: Be patient! The onboarding tutorial trains users incrementally--with repetition--to answer questions and assuage fears.

Onboarding Users (Meeting Demand)

Lesson Learned: “Enough” is never enough, convenience trumps cost, and users will pay for POSIX*.

Day 1: first group onto Stratus buys 20TB block storage (10% of usable)

- 1st of 31 known dbGaP projects
- Silver Lining:
  - Cost recovery ensures that we can scale to meet demand
  - If 500TB s3Cache is unused, capacity can be converted to block storage

* We are experimenting with Minio Client to make S3 interaction feel more like POSIX

Future Directions

Lesson Learned: dbGaP is just the beginning.

- HIPAA and FISMA (FedRAMP) are “desired”, but “need” is a lot harder to establish. (Build it and they will come?)

Lesson Learned: Give and inch and they’ll take a mile. General users desire a general compute cloud. Staff desire a new internal infra-cloud.

- Flexibility
- Price
- Security
- Support

Thank You!

Questions?

bollig@umn.edu

Additional

Organization

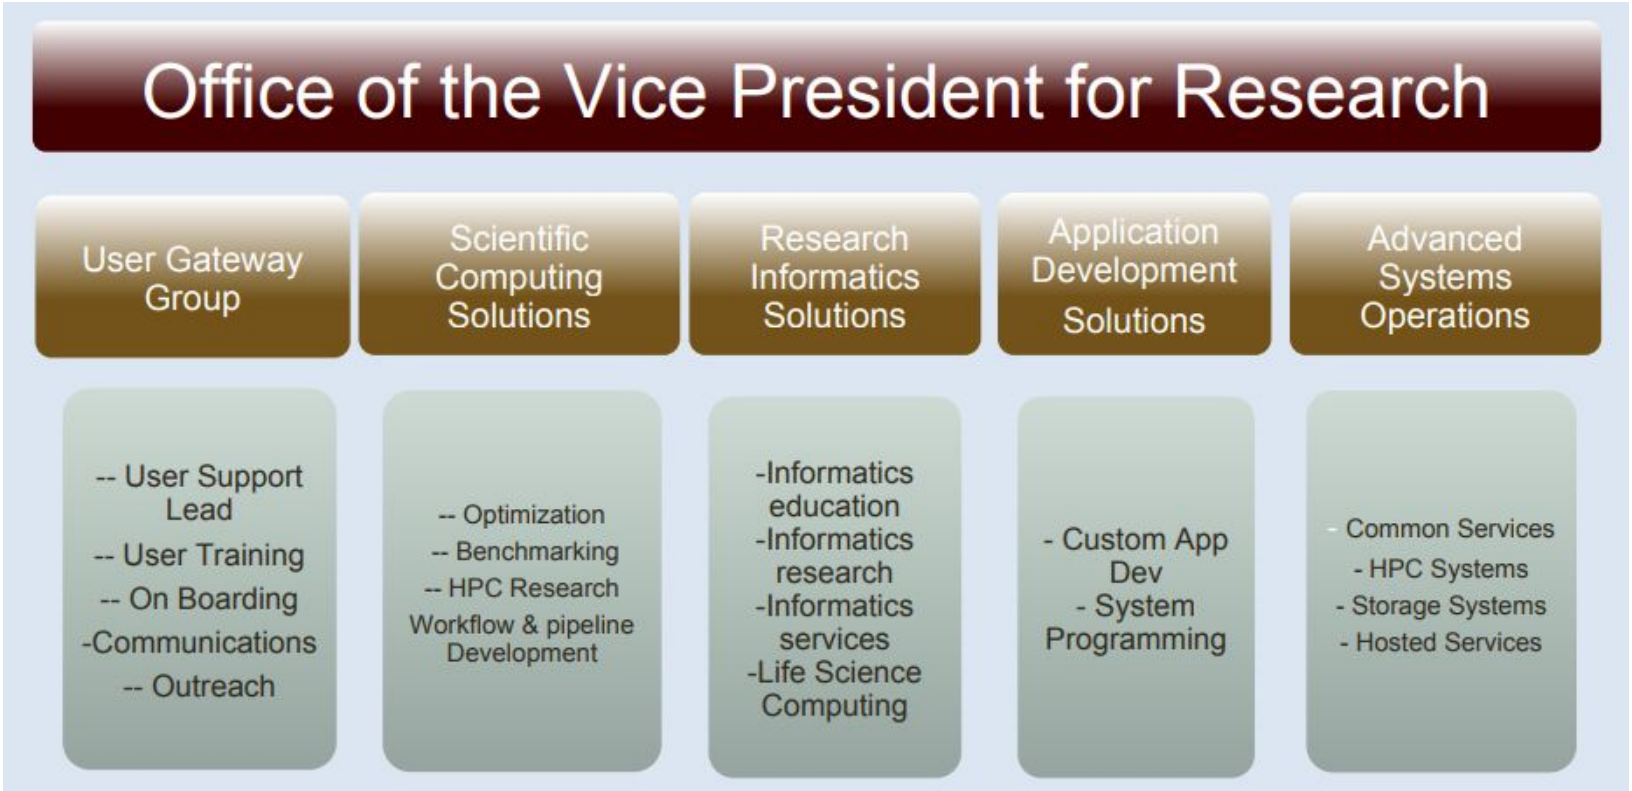

MSI Services

HPC Systems -- Two systems

- Mesabi (2015) still in Top 20 university-owned supercomputers in the nation (670 TFLOPs plus 105 TFLOPs from GPUs)

Storage -- Three Tiers

- 4.1 PB high-performance Tier I global storage
- 3.1 PB S3 Tier II Object Storage
- Tape Archive

Consulting

- Scientific Computing Solutions
- Research Informatic Solutions
- Application Development Services

Requirements and Planning (Other Motivations)

Subscription-based Service ensures funds exist to scale resources

Isolation from MSI core services improves integrity of all MSI services

Free and open source software with large community

In-house knowledge

Requirements and Planning

Service limits

- 30 day maintenance window
- Compulsory updates on single OS
- Global filesystem shared by all users
- One-off configurations to avoid storage backups and individualize ACLs
- Two-factor auth only through separate bastion

Virtualization Features

- Live-migration for long running jobs
- Self-service VMs updated independently to ensure compatibility
- Per-project tenants and user-lists
- Software defined networking with per-network rules
- Secure file storage