THE UNIVERSITY OF
CHICAGO
BIOLOGICAL
SCIENCES

Overview of CRI Computing Infrastructure

Mike Jarsulic
Olumide Kehinde

January 16, 2020

Introduction

IT Operations & Infrastructure Team

Scientific Computing Team

Three main service lines:

- Data Storage (Lab shares)
- Virtual Servers and System Administration
- HPC, Deep Learning, Commercial Application Servers

Procedures following the BSD's Information Security Policies to be able to store and process ePHI (HIPAA) data for research use

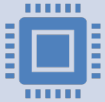

All primary resources are behind firewalls in the 6045 Kenwood datacenter

Data is backed up to tape in the 1155 (E 60th St.) datacenter

Staff

- Director
  - Thorbjorn Axelsson
- IT Operations and Infrastructure
  - Olumide Kehinde (Manager)
  - Joe Sutton (Sr. Windows Administrator)
  - Joel Van Os (Sr. Windows Administrator)
  - Adebimpe Akinlolu (Linux Administrator)
  - Jonathan Wroblewski (Linux Administrator)
  - Ellie Xu (Help Desk Analyst)
  - David Gonzalez-Velez (Help Desk Analyst)
- Scientific Computing
  - Mike Jarsulic (Manager)
  - Yi Du (HPC Administrator)
  - Camila Cuestas (Student)

Main Service Lines

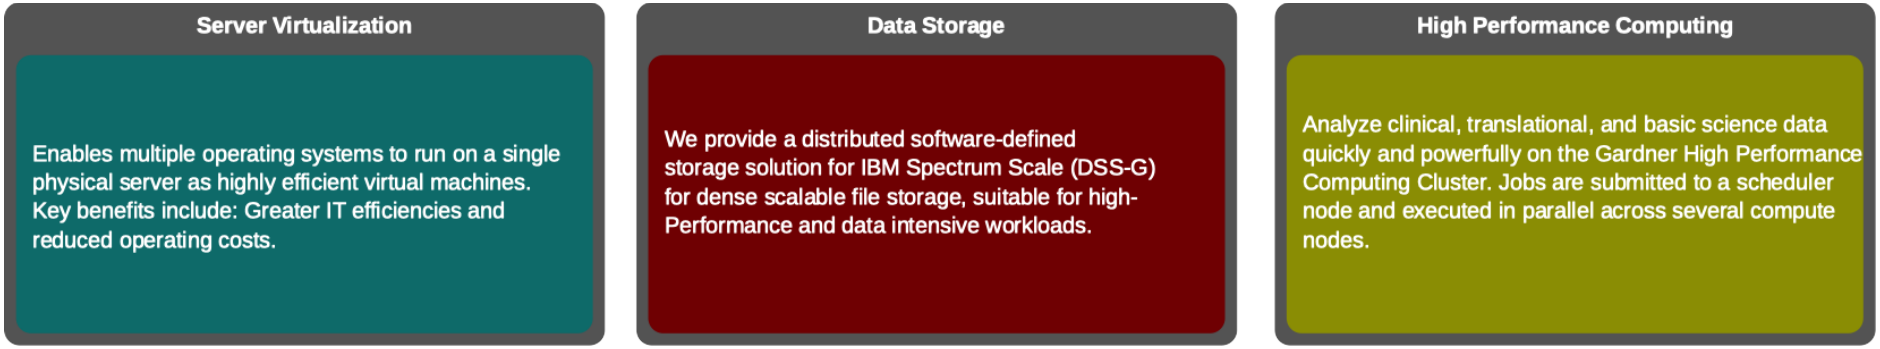

Others Services:

- Clinical Research Data Warehouse: contains clinical de-identified data dating back to 2006.
- Cohort Discovery (i2b2)
- Bioinformatics Core: Analysis of complex genomic data
- Dev Team: Development of custom solutions tailor-made for individual research projects.
- RedCap: A self-managed, secure, web-based solution designed to support data collection strategies for research studies

Virtual Environment

- Creates virtual resources that can be used the same way as any physical resource or application.
- Virtual resources include storage, CPUs, memory and networks.
- Allows deployment of custom applications.
- We currently support Windows 2012 R2, Windows 2016, RedHat Enterprise Linux 7 and 8.

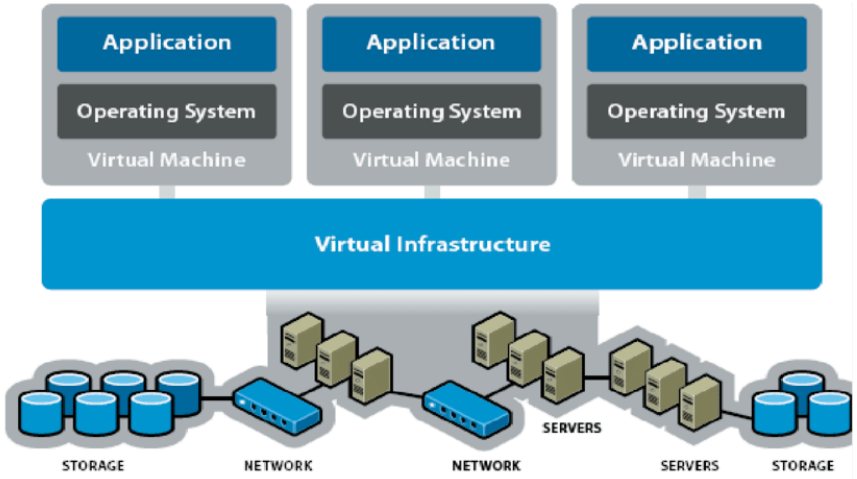

Virtual Environment

When you request a virtual machine, you will need to provide:

Purpose of the VM

Data Sensitivity: PHI, PII, PCI or No sensitive data (PHI: protected health information, PII: personally identifiable information, PCI: payment card information)

Expected Availability: 24/7, business hours, anytime

Environment: Test, Development, UAT, QA, Production

Network Access: Campus, Internal Only or External

Operating System: Windows or Linux

Server Configuration: Basic, Advanced or Enterprise

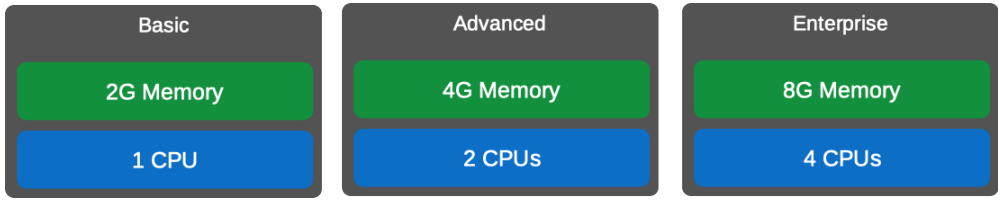

Application Services

- Applications we support include: Apache, IIS, MS SQL, MySQL, NGinx, Oracle

- **WinStats Server:** Is a Windows Server that hosts the following statistical applications - IBM SPSS, Statistics 24, Illumina GenomeStudio 2.0, SAS 9.4, Matlab, Jupyter Notebook, R Studio, Stata 14 and 15.

  Specifications:
  - 2 CPUs: 2.4GHz Intel Xeon E7-8870
  - 10 cores per CPU
  - 512 GB of RAM
  - 1.18 TB of dedicated SSD storage

- We also support two Linux stats servers (Stata, SAS, Matlab, Mathematica)

  Specifications:
  - 2 Intel Xeon E5-2698 v3 CPUs
  - 2.3 GHz
  - 16 ops/cycle
  - 16 cores per CPU
  - 768 GB RAM per server

Data Storage

Old Storage Systems

Problems with Isilon Cluster:

- High expansion cost
- High maintenance cost
- Ethernet requirement
- Not POSIX-compliant (parallel file access)
- Low read iops

Problems with Scratch Storage:

- Glusterfs
- Poor performance on small files
- Poor metadata performance
- Stability

New Storage System

Lenovo Distributed Storage Solution for IBM Spectrum Scale (DSS-G)

COMBINES SCRATCH AND PERMANENT STORAGE INTO ONE SYSTEM

LOWER COSTS

INCREASED PERFORMANCE

UTILIZES INFINIBAND NETWORK

POSIX-COMPLIANT

MAINTAINS STABILITY AND SECURITY OF THE ISILON

Data Storage (Cont'd)

Components:

2 Lenovo x3650 M5 controller nodes, 6 High density JBODs, 3 Protocol nodes, 1 Xcat node, 1 Backup node

Purpose:

Primarily for research data

Capacity:

- Shares - 3.1PB
- Home - 11TB (10GB quota)
- Scratch Space – 250TB (purged periodically)

Access Methods:

- SMB/CIFS from Windows, Mac and Linux clients
- Graphical SFTP clients: WinSCP, CoreFTP (Windows Only), FileZilla, Cyberduck (Windows & Mac)
- NFS access from CRI servers
- Native GPFS on HPC Cluster
- Rsync, sftp on cri-syncmon.cri.uchicago.edu

Data Backup

- File level backup for physical servers using IBM Spectrum Protect (formerly called Tivoli Storage Manager)
- Full Virtual Server backup using Spectrum Protect VE
- PRFS Storage is also backed up using a dedicated backup node using Spectrum Protect
- Backed up data gets migrated automatically from our Primary DC to a tape library at our DR site

Usage (Storage & VMs)

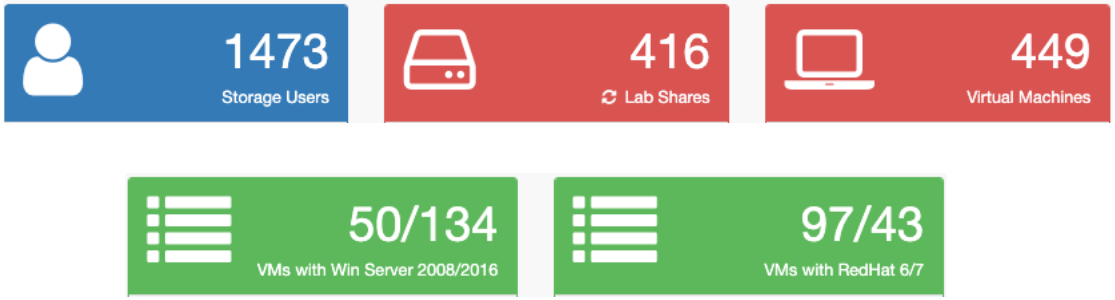

Departments Using CRI Storage

- Anesthesia And Critical Care
- Ben May Cancer Research
- Biochemistry and Molecular Biology
- BSD Academic Administration
- Center for Health and Social Sciences
- Committee on Evolutionary Biology
- Ecology and Evolution
- Family Medicine
- Grossman Institute for Neuroscience, Quantitative Biology and Human Behavior
- Human Genetics
- Medicine
- Microbiology
- Molecular Genetics and Cell Biology
- Neurobiology
- Neurology
- Obstetrics And Gynecology
- Office of Shared Research Facilities
- Ophthalmology and Visual Science
- Organismal Biology and Anatomy
- Orthopaedic Surgery and Rehabilitation Medicine
- Pathology
- Pediatrics
- Pritzker School of Medicine
- Psychiatry
- Psychology
- Public Health Sciences
- Radiation and Cellular Oncology
- Radiology
- Surgery

THE UNIVERSITY OF CHICAGO MEDICINE & BIOLOGICAL SCIENCES
Center for Research Informatics

Resource Access

In order to access CRI resources:

- You need to have an active BSD account
- You need to request resource access using our Web Provisioning Forms (Link available on the CRI website)
- VPN Access (Campus or BSD) may be required for some resources either on or outside campus.

After submitting the request:

- A ticket will be created, assigned and processed.
- You can follow up on the status of the ticket by replying to the automated ticket response email.
- You will be sent access instructions.

Requesting Access to CRI Resources

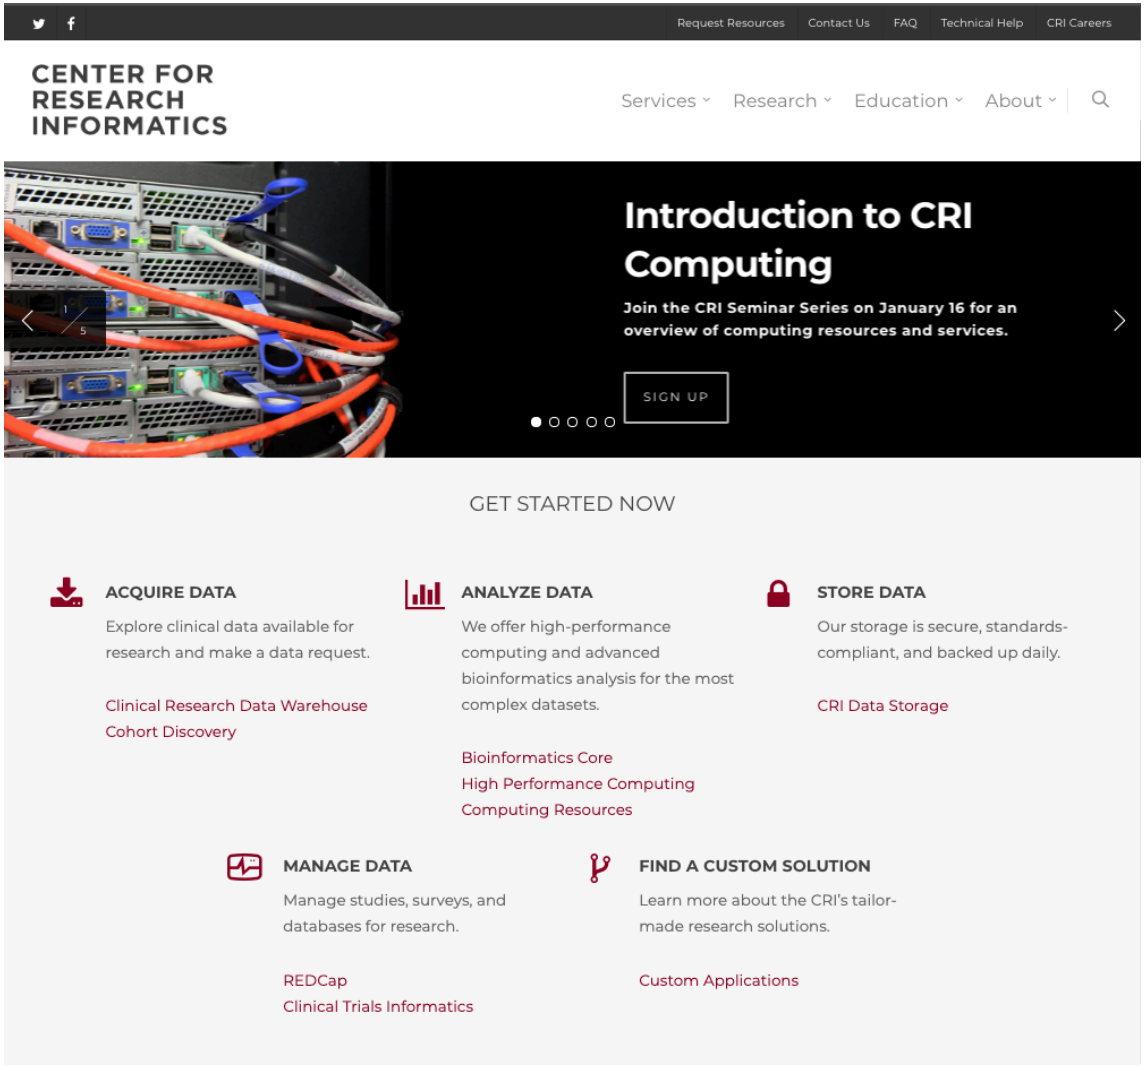

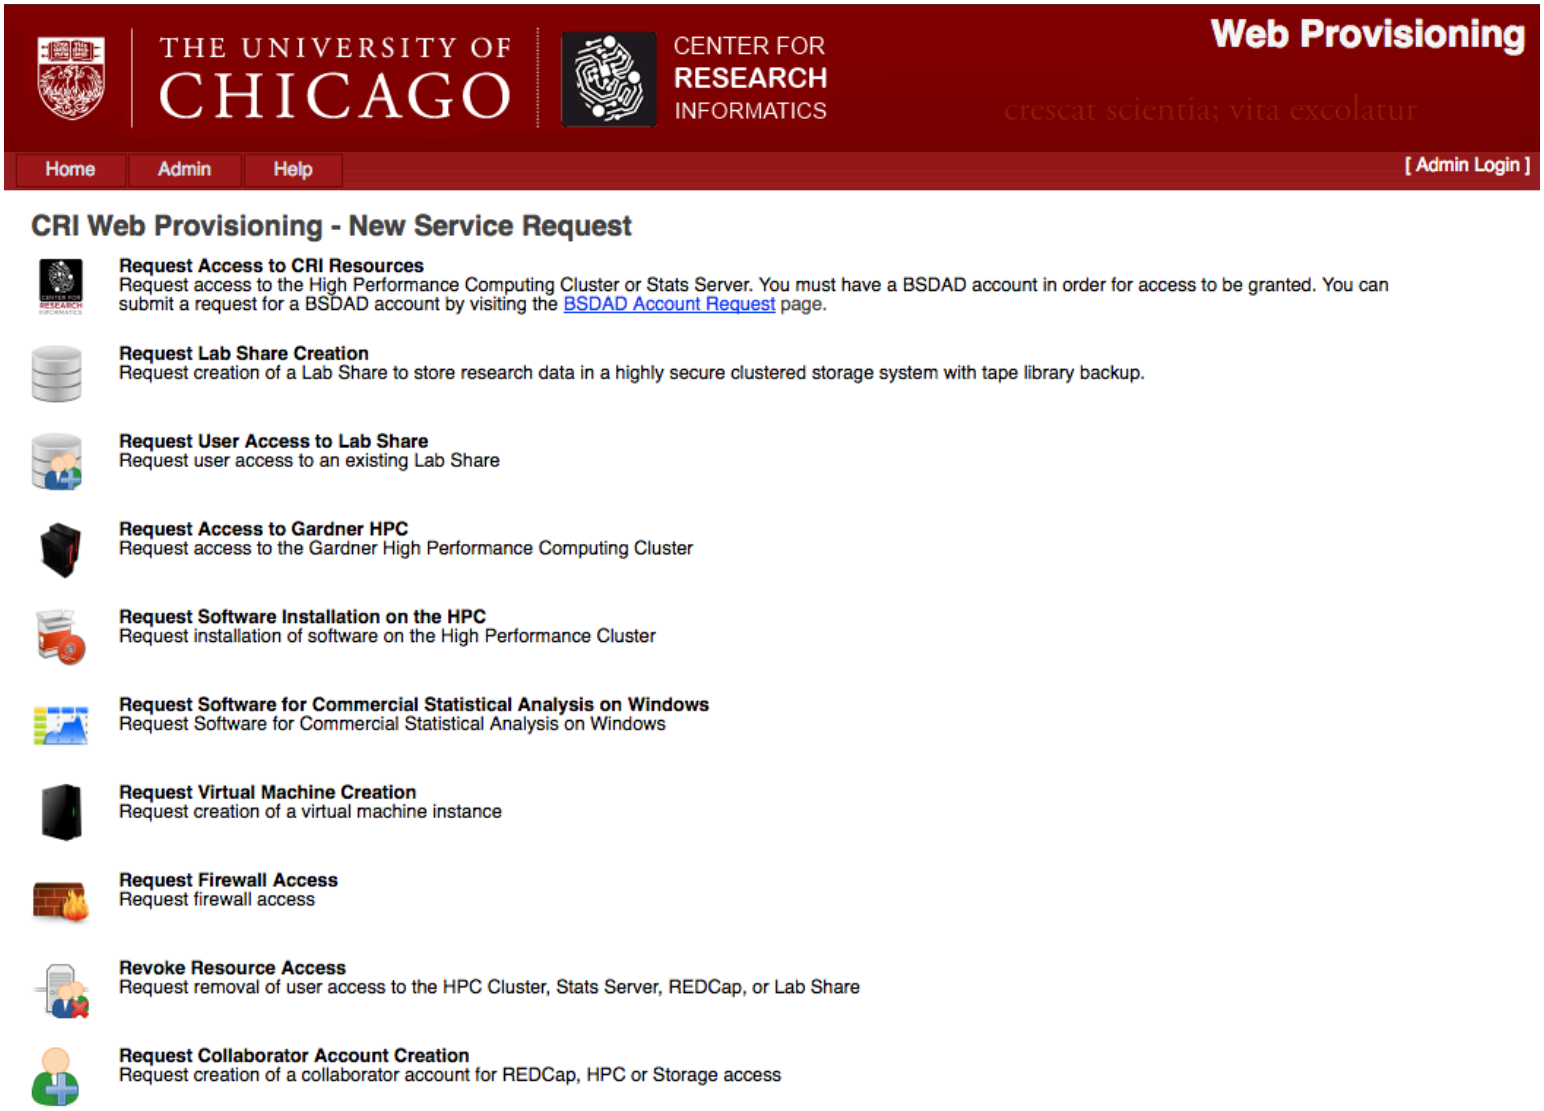
THE UNIVERSITY OF CHICAGO
CENTER FOR RESEARCH INFORMATICS
Web Provisioning
crescat scientia; vita excolatur
Home
Admin
Help
[Admin Login]
CRI Web Provisioning - New Service Request
Request Access to CRI Resources
Request access to the High Performance Computing Cluster or Stats Server. You must have a BSDAD account in order for access to be granted. You can submit a request for a BSDAD account by visiting the BSDAD Account Request page.
Request Lab Share Creation
Request creation of a Lab Share to store research data in a highly secure clustered storage system with tape library backup.
Request User Access to Lab Share
Request user access to an existing Lab Share
Request Access to Gardner HPC
Request access to the Gardner High Performance Computing Cluster
Request Software Installation on the HPC
Request installation of software on the High Performance Cluster
Request Software for Commercial Statistical Analysis on Windows
Request Software for Commercial Statistical Analysis on Windows
Request Virtual Machine Creation
Request creation of a virtual machine instance
Request Firewall Access
Request firewall access
Revoke Resource Access
Request removal of user access to the HPC Cluster, Stats Server, REDCap, or Lab Share
Request Collaborator Account Creation
Request creation of a collaborator account for REDCap, HPC or Storage access

Getting Help

Contact the CRI using our Support Portal:

- cri.uchicago.edu > Technical Help >
- Fill out support form
- This will create a support ticket
- Filling out the form ensures we collect information necessary to start processing the ticket

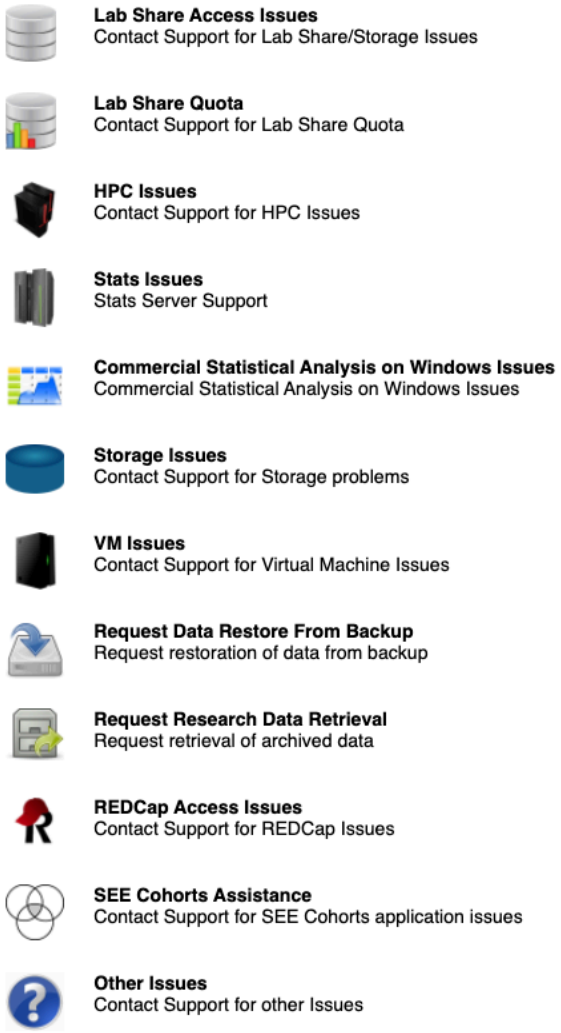

2020 Goals

Server Migrations/Upgrades:

- Complete BSD Active Directory Domain Upgrade
- Windows 2008R2 to 2016 Server Migration
- Windows 7 to Windows 10
- RHEL 6 to RHEL 8 (or RHEL 7) Migration.

Spectrum Scale Test Environment Buildout

Spectrum Scale Upgrade:

- Allows continued support of native GPFS protocol access in the upgraded HPC environment
- Increase efficiency of parallel workloads
- Use of sub-blocks in the future for efficient management of small files
- Increase in metadata operations efficiency
- New code supports built-in file auditing capabilities (license upgrade required)
- Maintain the same level of system stability, compliance, and data integrity

2020 Goals (Cont'd)

- Redeploy our Internal Logging Cluster using Windows Collector & ELK Stack
- Upgrade our monitoring system (Nagios XI)
- Upgrade VMWare
- Replace aging firewall appliances
- Replace aging networking devices
- Upgrade IPA (RedHat Identity Management) Cluster
- Setup VEEAM Backup
- Migrate Workstations to Asset Panda

High Performance Computing – What is a HPC Cluster?

- Storage
- Login Nodes
- Scheduler Node
- Compute Nodes
- GPUs
- Xeon Phi

High Performance Computing – Storage

- PRFS Storage Cluster
  - Home Directories (/home/<userid>)
  - Lab Shares (/gpfs/data/<labid>)
  - Applications (/apps/software)
  - Total Size: 3.1 PB
  - Bandwidth: 56 Gb/s
  - Permanent, Quota'd, Backed Up,
  - Available everywhere

- Scratch Space
  - /scratch/<userid>
  - Bandwidth: 56 Gb/s
  - Total Size: 250 TB
  - Purged, Private, Not Quota'd, Not Backed Up
  - Limited availability

High Performance Computing – Login Nodes

- Purpose
  - User interface for the HPC cluster
  - Composing editing jobs
  - Submitting jobs
  - Tracking/Managing jobs
  - Writing source code
  - Compiling source code

- DO NOT RUN ANALYSIS ON THE LOGIN NODES!!!

High Performance Computing – Scheduler Node

- Purpose
  - Keeps track of which resources are available on each compute node of the cluster
  - Schedules jobs based on available resources
  - Maintains historical metrics on jobs

High Performance Computing – Compute Nodes

- Specifications
  - Nodes:
    - 88 standard
    - 28 mid-tier
    - 4 high-tier
  - Processors (2 per node):
    - Intel Xeon E5-2683 v3
    - 14 cores
    - 2.0 GHz
    - 16 ops/cycle
  - Memory
    - Standard: 4 GB/core
    - Mid-tier: 16 GB/core
    - High-tier: 45 GB/core

High Performance Computing – Accelerator Nodes

- GPUs
  - 5 Nodes
  - 1 NVidia Tesla K80 per node
    - Contains 2 NVidia Tesla GK210 GPUs
    - 2496 CUDA cores
    - 24 GB memory
- Xeon Phi
  - 1 node
  - 2 Intel Xeon Phi P5110 coprocessors
    - 60 Cores per coprocessor
    - 8 GB RAM

High Performance Computing – Software

- Compilers
  - GNU
  - Intel
  - PGI
  - Java JDK
- Scripting Languages
  - Perl
  - Python
  - R
- Software Environment
  - LMod

High Performance Computing – Obtaining an Account

- Prerequisites: BSD Account
- Sign up for and account
  - http://cri.uchicago.edu
  - Experience Level
  - Software Requests
  - Email Address for Job Output
  - Emergency Phone Number
- Collaborator Accounts

HPC Metrics

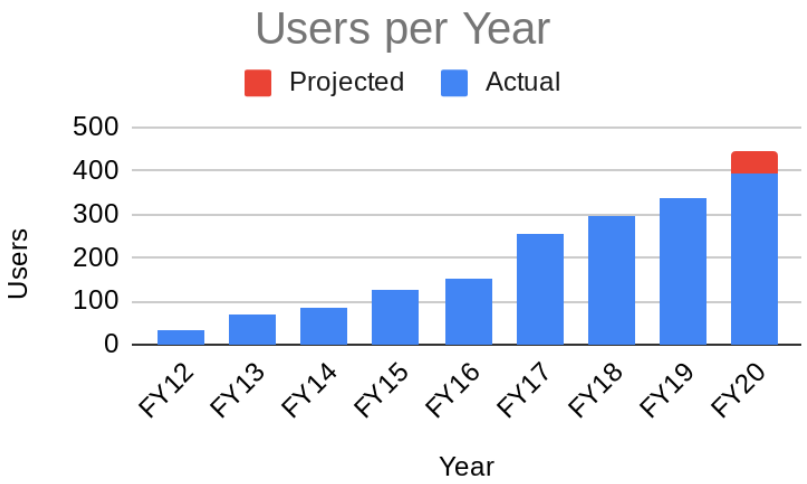
Users per Year
Projected
Actual
500
400
300
200
100
0
Users
FY12
FY13
FY14
FY15
FY16
FY17
FY18
FY19
FY20
Year

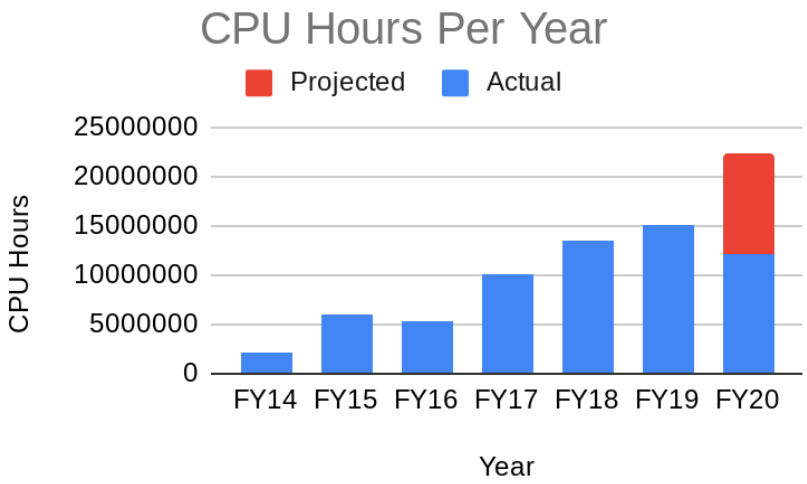
CPU Hours Per Year
Projected
Actual
25000000
20000000
15000000
10000000
5000000
0
CPU Hours
FY14
FY15
FY16
FY17
FY18
FY19
FY20
Year

Deep Learning

- 2 Nodes
  - 2 Intel Xeon Gold 6230 CPUs
    - 20 cores per CPU
    - 2.1 GHz
  - 768 GB memory
  - 8 NVidia Tesla V100 per node
    - 5120 CUDA Cores
    - 640 Tensor Cores
    - 16 GB memory
    - Double Precision: 7 TFLOPs
    - Single Precision: 14 TFLOPs
    - Tensor: 112 TFLOPs

Scientific Computing - Support

- How to get help?
  - Email: hpc@rt.cri.uchicago.edu
  - Support form on the CRI Website
  - Documentation: Coming Soon
  - Ask a friend
  - Office Hours (Tuesday/Thursday)
  - Bioinformatics Core
  - CRI Seminar Series
- How not to get help?
  - Calling me (unless it's an emergency)
  - Email me directly

Future: HPC Upgrade

- Red Hat Linux Enterprise 7.6
- SLURM (scheduling/resource management)
- Container capabilities
- Spack
- LMOD metrics
- Compilers
  - gcc 4.9.4
  - gcc 8
  - gcc 9
  - Intel 2020
  - PGI 2019
  - llvm

PBS vs. SLURM

#PBS –N job1	#SBATCH --job-name=job2
#PBS –l nodes=1:ppn=4	#SBATCH --nodes=1
	#SBATCH --ntasks=4
#PBS –l walltime=6:00:00	#SBATCH --time=6:00:00
#PBS –l mem=4gb	#SBATCH --mem=4096
#PBS –o stdout.log	#SBATCH --output=stdout.log
#PBS –e stderr.log	#SBATCH --error=stderr.log

Future Projects

- New HPC Cluster
- Archive Storage
- Virtual Infrastructure Upgrade
- Logging Infrastructure Replacement
- Winstats Replacement
- Stats Replacement

Thank you!

Questions?