



BRC HPC Services/Savio

Krishna Muriki and Gregory Kurtzer

LBNL/BRC

kmuriki@berkeley.edu, gmk@lbl.gov

Berkeley
UNIVERSITY OF CALIFORNIA

SAVIO - The Need Has Been Stated

Inception and design was based on a specific need articulated by Eliot Quataert and nine other faculty:

- *Dear Graham,*

- *We are writing to propose that UC Berkeley adopt a **condominium computing model**, i.e., a more centralized model for supporting research computing on campus...*

Review of May 2015 Town Hall

Faculty Computing Allowance was announced.

Topics Covered

- BRC Program Overview
 - Savio cluster technical architecture
 - Savio user environment and user support

Presentation slides and Video recordings available

- Online [here](#)
- Search @ research-it.berkeley.edu

SAVIO - Condo Service Offering

- Purchase into Savio by contributing standardized compute hardware
- An alternative for running a cluster in a closet with grad students and postdocs
- The condo trade-off:
 - Idle resources are made available to others
 - There are no (ZERO) operational costs for administration, colocation, base storage, optimized networking and access methods, and user services
- Scheduler gives priority access to resources equivalent to the hardware contribution

Faculty Computing Allowance

- Provides allocations to run on Savio as well as support to researchers who have not purchased Condo nodes
- Up to 200k Service Units (core hours) annually
- More than just compute:
 - File systems
 - Training/support
 - User services

Faculty Computing Allowance (Cont.)

- Eligibility requirements
 - UCB campus Faculty
 - Allowance can be used by the faculty and/or by immediate group members.
 - In need of compute power to solve a research problem.
 - Meets the readiness requirements
- Allowance Request Procedure
 - First fill out the Online Requirements Survey
 - For additional cluster accounts: - Additional User Account Request Form
- Allowances
 - New allowances start on June 1st of every year.
 - Mid-year requests are granted a prorated allocation
 - A cluster specific project (*fc_projectname*) with all user accounts is setup
 - Scheduler account (*fc_projectname*) with the allocation is setup
 - Annual allocation expires on May 31st of the following year

SAVIO - Usage Readiness

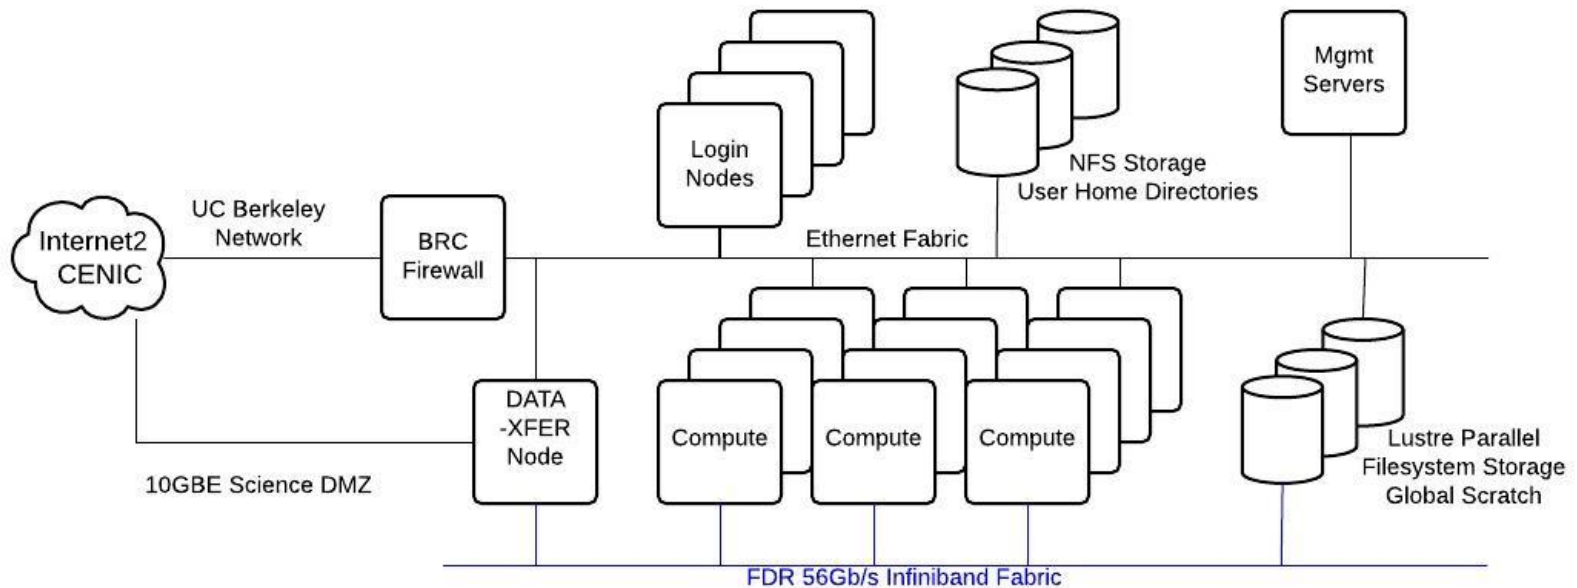
- Applications that can run on RHEL 6.
- Familiarity with user application stack usage and maintenance on Linux OS
- Ability to use a shared resource in a fair manner
- No requirement of application specific persistent database hosting within cluster infrastructure.

Stats - Allowances, Condos & Users

- Faculty Computing Allowance
 - Started on June 1st 2015
 - Participated faculty count - 43
 - Number of groups setup - 40
- Condo groups
 - Number of faculty participated – 17
 - Number of Condo groups - 11
- Number of unique user accounts - 248

SAVIO - System Architecture

SAVIO System Overview



SAVIO - System Overview

- Similar in design to a typical research cluster
 - Master Node role has been broken out (management, scheduling, logins, file system, etc..)
- Home storage: Enterprise level, backups, 10GB
- Scratch space: Large and fast (~1PB Lustre)
- Multiple login/interactive nodes
- DTN: Data Transfer Node
- Compute nodes are delineated based on role

SAVIO - Online Resources

- Online User Documentation
 - research-it.berkeley.edu
 - [User Guide – Available online](#)
 - [New User Information - Available online](#)

One-Time Passwords (OTP)

- The biggest security threat that we encounter ...

STOLEN CREDENTIALS

- Credentials are stolen via keyboard sniffers installed on researchers laptops or workstations, incorrectly assumed to be secure
- OTP (One Time Passwords) offers mitigation
- Easy to learn, simple to use, and works on both computers and smartphones!

SAVIO - Specifications

- First Generation Hardware
 - Compute Nodes: 20-core, 64GB, InfiniBand
 - BigMem Nodes: 20-core, 512GB, InfiniBand
- Software Stack
 - Scientific Linux 6 (equivalent to Red Hat Enterprise Linux 6)
 - Parallelization: OpenMPI, OpenMP, POSIX threads
 - Intel Compiler
 - SLURM job scheduler
 - Software Environment Modules

SAVIO-2 Expansion

- Second Generation Hardware
 - Compute Nodes: 24-core, 64GB, InfiniBand
 - Nodes with Intel Haswell processors
 - 72 nodes Institutional contribution.
 - Current hardware specs for condo contributions.
- Savio2 nodes with GPUs
 - GPUs are optimal for massively parallel algorithms
 - 4 Kepler GPU nodes with support for ECC
 - 6 Specialized nodes (64GB, dual Nvidia K80)
 - 24 total GPUs

SAVIO - Pilot experiments

Based on user feedback & demand

- Serial/HTC Jobs
 - Support for less than node allocations
 - Nodes are shared between jobs
 - Specialized node hardware (12-core, 128GB, PCI flash storage, faster clock rate)
 - SSD/Flash based node local storage option

SAVIO - Pilot experiments (cont.)

- Low Priority Preemption queues
 - Special partition – jobs get low priority, can be killed.
 - Allows users in need to scavenge unused cycles
 - Made available to few early testers
 - Will soon be expanded to all the condo users.

SAVIO - How to Get Help

- Helpdesk
 - Email : brc-hpc-help@lists.berkeley.edu
 - Monday - Friday, 9:00 am to 5:00 pm
 - Best effort in non-working hours

Thanks! Questions?

For more information visit:

research-it.berkeley.edu/brc/hpc

Email:

research-it@berkeley.edu