

Virtual Clusters in the Jetstream Cloud

A story of Elasticized HPC

July 29th, 2019

Humans in the Loop: Enabling and Facilitating Research on Cloud Computing
PEARC 19 Chicago, IL

Eric Coulter

XSEDE Cyberinfrastructure Resource Integration
Indiana Univ. Science Gateways Research Center

XSEDE

Extreme Science and Engineering
Discovery Environment



Supported by OAC 15-48562.

Background

- XSEDE CRI (Formerly Campus Bridging) + IU SGRC (Science Gateways Research Center)
- Strive to make researchers' use of computational resources feel as simple as running calculations on their own laptops
- Enable sysadmins and local resource providers to build XSEDE-Similar resources
- Gateways provide portals for seamless use of **many** resources across a scientific domain
- How can we effectively leverage cloud resources for highly parallel codes?



XSEDE

VC Design

- Let's do the “wrong” thing!
 - Stick a traditional HPC System in the cloud!
- Designed to be accessed by Science Gateway users
- KISS
 - single headnode/login node – persistent IP
 - single internal network
 - shared NFS filesystem
 - heterogeneous compute nodes
 - single user for job submissions



VC Ecosystem

- Jetstream Cloud
 - Openstack-based research cloud
 - Hosted at IU and TACC
 - Available as an XSEDE Resource
 - So, able to draw interested users from ECSS projects, Gateways, CC, etc.
- Cloud API is well contained; able to abstract to other commercial cloud providers, though recently cluster offerings have started appearing commercially...



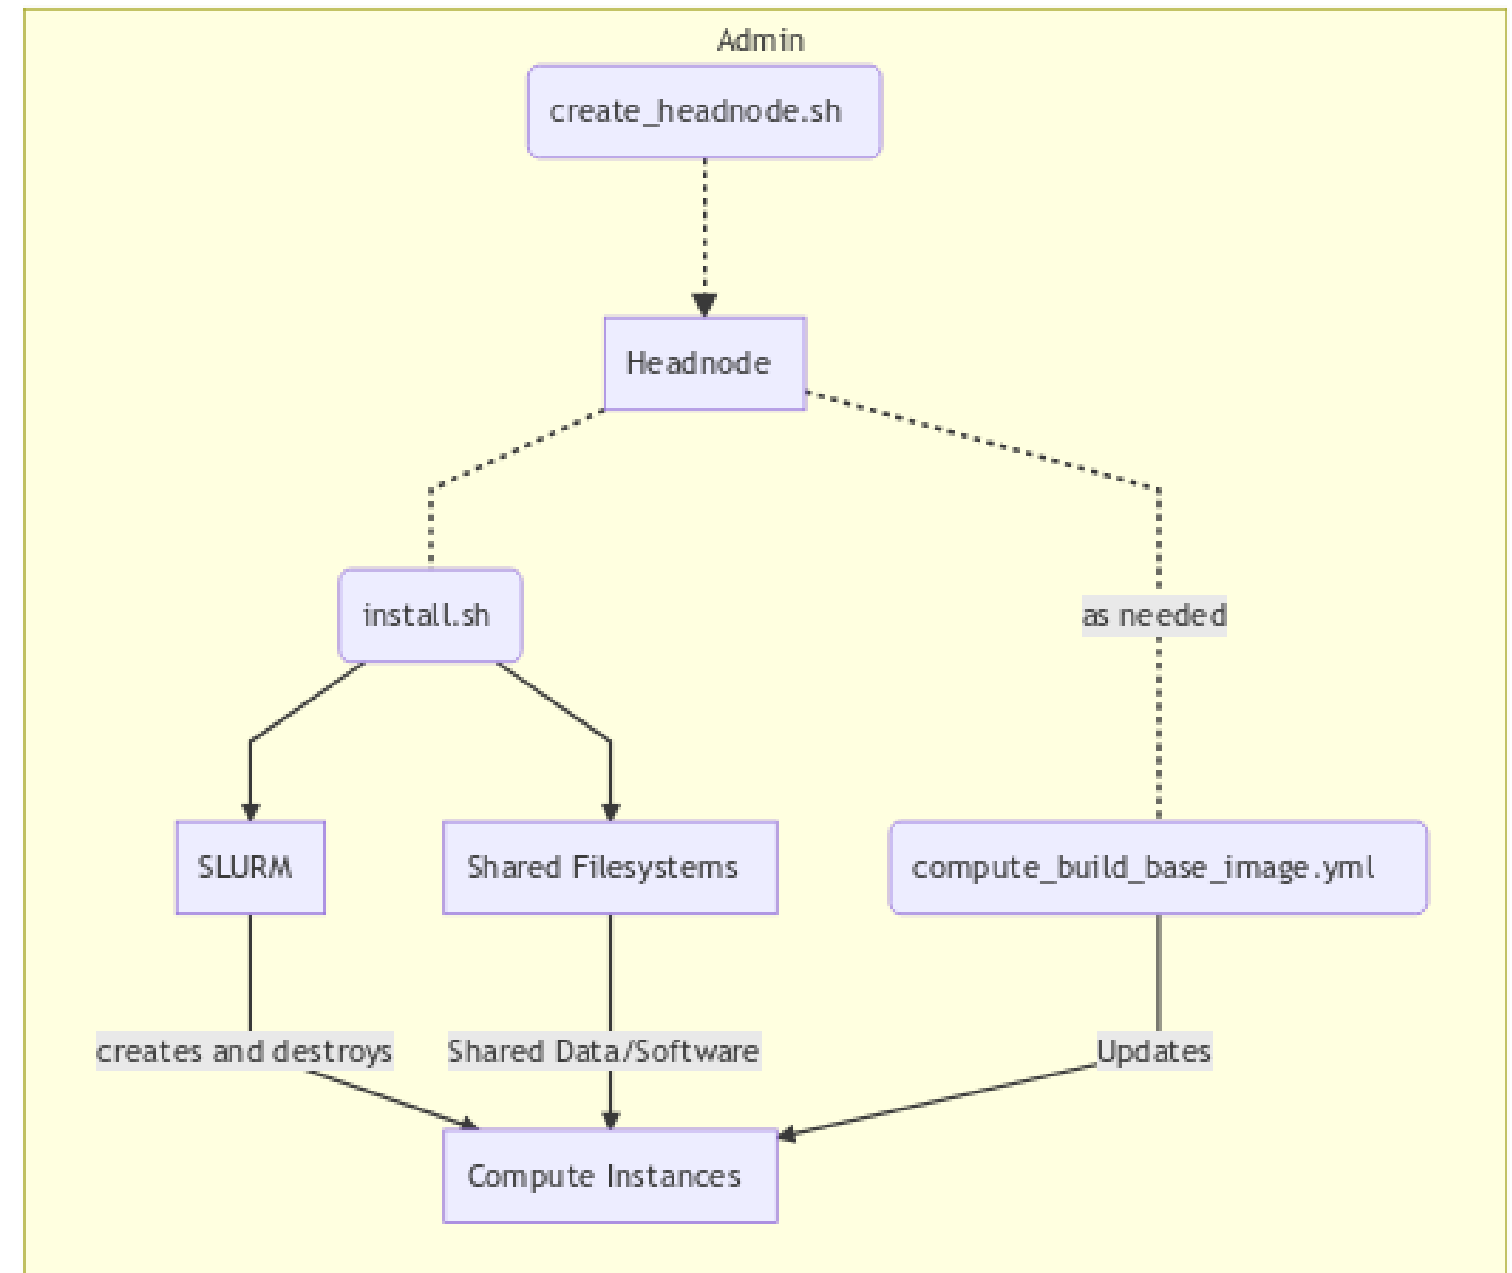
VC Evolution

- Let's continue doing the wrong thing:
- Part 1: Static compute instances (Maintenance nightmare)
- Part 2: Completely rebuild compute instances **EVERY SINGLE TIME THEY REBOOT (SLOW)**
 - Don't destroy the compute instances – just suspend/resume (Still a maintenance nightmare!)
- Part 3: Save an image of the compute nodes, only rebuild as needed, completely destroy nodes as needed.
 - Much less wrong!



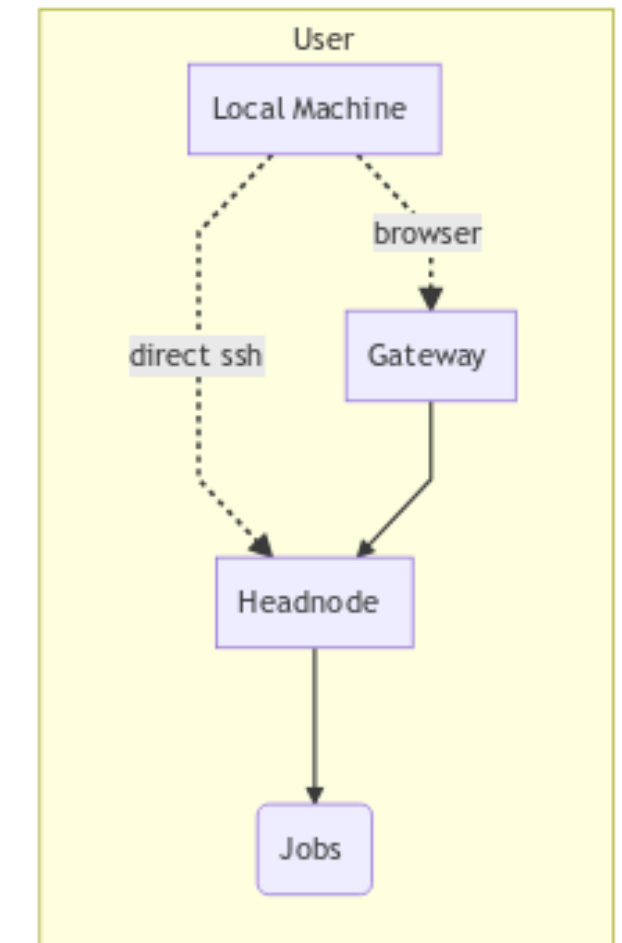
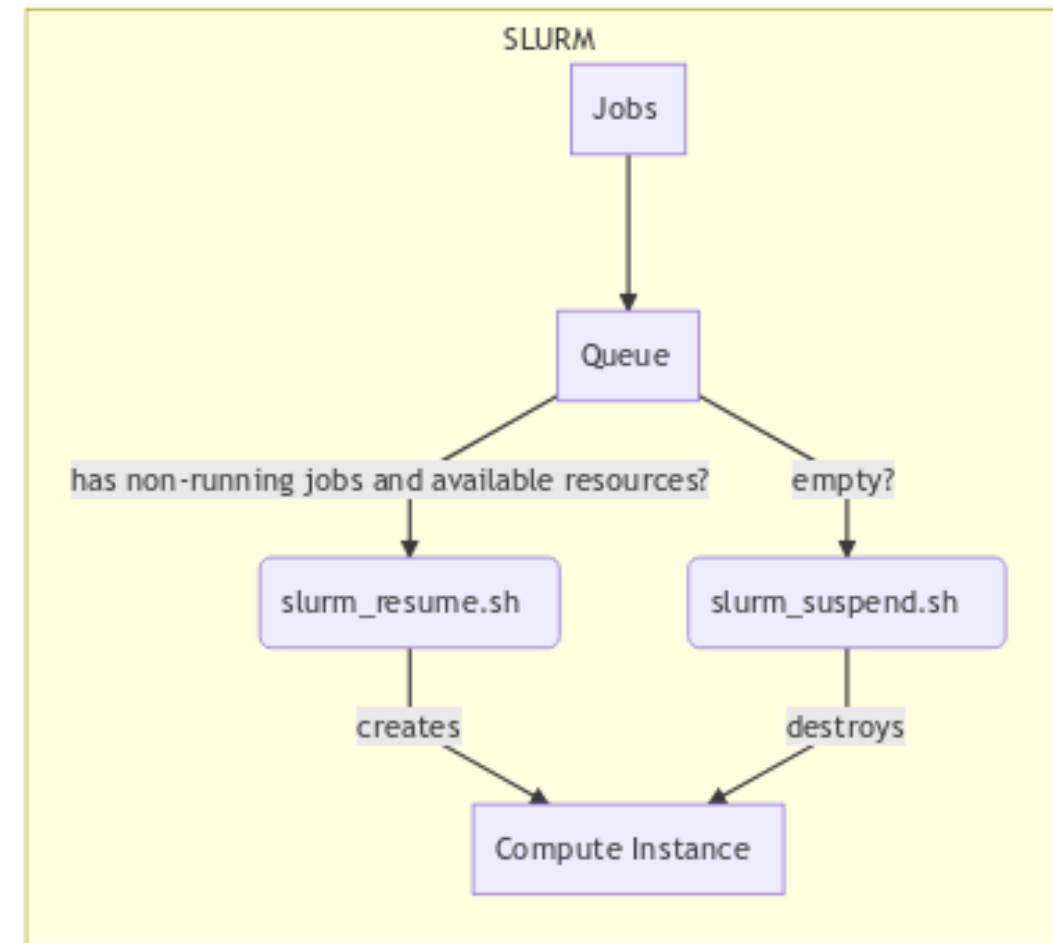
VC Architecture

- Let SLURM do the lifting
- Attempt to minimize human intervention
- Front-load effort into initial build and configuration
- Aim for fast, reliable job completion



VC User Interaction

- Standardize on Gateway users
- Allow flexibility for interactive use as well
- No need to deal with creation/destruction of cloud infra once the headnode is up



Education

- As an offshoot, this project has spawned a series of tutorials on cluster building aimed at HPC users or new admins
- Doing things the hard way, before playing with automated versions
 - Long division before short, plum-pudding before electron clouds, etc.
- Consistent good reviews; similar in spirit to our on-site work
- Cloud environment makes a powerful platform for sysadmin education – without neglecting to include modern principles of cloud management!
- Some level of understanding of low-level details still necessary at this stage




Projects

- ~10 VCs currently in production
 - there could be more!
- Lessons learned during this also applied to Galaxy Infrastructure
- Managed by 0.25 FTE
- Primarily interacted with via Airavata-powered Gateways
- Everything from MD to Textual Analysis



Projects: SEAGrid

- The Science and Engineering Applications Grid
- Offers access to Computational Chemistry, Fluid Dynamics, Structural Mechanics, and Fluid Dynamics applications
- Spans a wide range of HPC systems – Big Red II, Comet, Stampede, and the first JS Virtual Cluster, with data shared from Wrangler



Projects: InterACTWEL

- Interactive Adaptation and Collaboration Tool for managing Water, Energy and Land
- PI-s at Oregon State and Indiana Univ.-Purdue Univ. Indianapolis
- Computer-aided decision tool for long-term planning at the Food-Energy-Water nexus
- Allows coordination and data sharing by stakeholders at all levels




XSEDE

Maintenance and Updates

- These things stick around for multiple-year projects!
- Headnode should remain available at all times
- Software requires updates!
 - (Security vulnerabilities at the very least)
- Often updates to scientific software are needed as well,
- which means new compute images!
- Occasional cloud hiccups require jumping through \$N VC systems, rebooting/fixing/testing.
- Software updates may be non-trivial!



Humans in the Loop



- Approaching the dream of hands-off clusters
- Software and funding peculiarities dictate a certain level of human effort
- Intersections of “human” problems require breaking automation

- Questions?

- Email: help@xsede.org for CRI Help
 - With XCRI in the subject
- : sgrc-iu-grp@iu.edu for Gateways
- Explore: github.com/XSEDE/CRI_JS_Cluster
(requires a Jetstream allocation + bash shell)

- Thanks to:



XSEDE

Extreme Science and Engineering
Discovery Environment



Supported by OAC 15-48562.