Introduction to Advanced Research Computing

Sarah Ghazanfari

Computational Scientist

Advanced Research Computing (ARC), Information Technology Division

Virginia Tech University

[Sign In]

VIRGINIA TECH

Summer 2024

ARC workshop Series, Summer 2024 [Sign In]

- 06/03 : Introduction to Advanced Research Computing

  Basics of HPC, computer clusters, HPC resources, access to ARC systems

- 06/03 : Connect to ARC Systems and Run your first jobs

  Connect via Open OnDemand, connect via SSH, cluster and scheduler orientation, run demo jobs

- 06/04 : Running code/software on ARC systems in different ways

  Job environments (modules and Conda), running interactive and batch jobs

- 06/04 : Launching Jobs in Parallel on ARC Clusters

  MPIRUN vs. SRUN, GNU parallel for load balancing, SRUN for resource detection and binding, "Built-in" or library-based parallelism

- 06/05 : Monitoring Resource Utilization and Job Efficiency

  Acquiring resources, characteristics of compute nodes, overall activity, current loads, job status

Outline

- Learning goals:
  - ✓ ARC structure
  - ✓ Mission
  - ✓ Computing resources
  - ✓ Where/how to get help
- Mostly informational about ARC and research computing at VT.
- We want to hear your questions.
- Feedback needed to help improve future workshops.

Sign In

VIRGINIA TECH

Scenarios bringing people to ARC

- **Scaling out:** “Our first analysis took four hours on a laptop. The results are great, but we need to run it 8,500 more times in a few months.”
- **Uninterrupted time:** “My jobs take 6 days on my office computer, but if it reboots due to an update, I lose everything.”
- **Scaling up:** “I need to process an 80GB dataset with a colleague’s program. My computer handles 3GB datasets fine, but crashes with the larger one. I think I need more memory.”
- **Platform for novel technologies:** “I want to test a neural network for my problem but training it on my data takes weeks.”

ARC Structure

Director and VP,
David Raymond

- Computational Scientists “Forward Team”
- Visualization
- Systems Engineering “Backend Team”

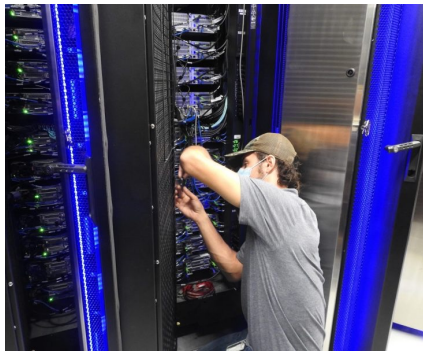

VIRGINIA TECH

The ARC Team

- **VP for Research Computing:** David Raymond
- **Visualization:** Nicholas Polys, Director
  Ben Sandbrook
- **Network Research Manager:** Mark K. Gardner
- **Computational Scientist:** Matthew Brown, Ayat Mohammed,
  Chris Kuhlman, Sarah Ghazanfari, Sofia Lima
- **Systems Engineering/Administration/Development:** Jessie Bowman, Miles Gentry, Jeremy Johnson,
  Nathan Liles, William Marmagas, Doug McMaster
- **Our student interns and Helpdesk GRAs:** Saikat Dey, Eslam Hussein, Sonal Jha, Samira Mali

https://arc.vt.edu/about/our-team.html

VIRGINIA TECH

Research Examples

- 1000+ of registered user accounts
- 400+ active projects each year
- 200+ publications (that we know about)

Geosciences, Economics, Mechanical Eng., Agriculture and Applied Economics, Aerospace and Ocean Eng., Computer Science, Entomology, Statistics, Civil and Environmental Eng., Industrial and Systems Engineering, Biomedical Eng., Plant Science, Physics, Forestry, Psychology, Accounting, Business Admin., Finance, Marketing, … more!

- ✓ ***"… experimentally measure the 3D Rayleigh index, which quantifies whether a combustion system is thermoacoustically unstable…"***
- ✓ ***"Perform large-scale computer simulations to recreate the sensory world of bats… to develop efficient sensing paradigms that are parsimonious yet suitable for complex, unstructured natural environments such as dense forests"***
- ✓ ***"… parallel computation of simulated structural components and systems subjected to mechanical loadings or chemical deterioration mechanisms"***

Research Examples

- 1000+ of registered user accounts
- 400+ active projects each year
- 200+ publications (that we know about)

Geosciences, Economics, Mechanical Eng., Agriculture and Applied Economics, Aerospace and Ocean Eng., Computer Science, Entomology, Statistics, Civil and Environmental Eng., Industrial and Systems Engineering, Biomedical Eng., Plant Science, Physics, Forestry, Psychology, Accounting, Business Admin., Finance, Marketing, … more!

- ✓ ***"…estimate hydrodynamic forces … in design, analysis, and optimization of swimming microrobots"***
- ✓ ***" genome assembly for the wild chili, Capsicum chacoense"***
- ✓ ***"teach students computational methods that scientists use to understand the brain at the anatomical level in order to gain insights into structure-function relations, health, and disease"***
- ✓ ***"… a dramatic increase in earthquake activity is a result of deep underground disposal of oilfield wastewater … understand the mechanisms driving fluid migration to seismogenic depths…"***

VIRGINIA TECH

ARC Services and Resources

Topics Overview:

- Mission and goals
- Resources and services
  - High Performance Computing / High Throughput Computing / Research Computing / Visualization
  - Consultation / Collaboration / Helpdesk
  - Teaching / Workshops / Instruction
- Getting started
  - Accounts / Accounting / Planning / Lifecycle
  - Walkthrough
- Getting assistance
  - Websites / Helpdesk / Office Hours / Consultation

ARC online documentation here:
https://www.docs.arc.vt.edu/get_started.html

We want you to use ARC resources. These resources can help you get your work done. How can we help?

VIRGINIA TECH

ARC's Mission

- Advanced Research Computing (ARC) provides ***centralized support*** for research computing by ***building, operating and promoting*** the use of advanced cyberinfrastructure at Virginia Tech.

- ARC delivers a ***comprehensive ecosystem*** consisting of advanced computational systems, large-scale data storage, visualization facilities, software, and consulting services.

- ARC provides ***education and outreach services*** through conferences, seminars, and scientific computing courses.

- ARC seeks to help ***maximize research productivity*** at Virginia Tech through interdisciplinary ***collaborations*** that connect researchers to new opportunities in computing and data driven research as they occur.

- ARC can become ***part of your team*** or ***support your team***: proposal writing (technical content, facilities descriptions, data security), software and workflow design, funding for prioritized system usage (beyond the free accounts).

- By ***fostering strategic partnerships*** with the public and private sector, ARC serves to cultivate an entrepreneurial spirit around advanced computing infrastructure as a platform for collaboration and helps secure the position of Virginia Tech as a leader in education and research.

VIRGINIA TECH

Resources & Services

- ✓ High performance computing
- ✓ High throughput computing
- ✓ Large memory computing
- ✓ Research computing
- ✓ Reliable/available computing
- ✓ Visualization

Note: your computing needs do -NOT- have to be "high anything" as in HPC. You may just want to get your computations off of your tower/laptop to make it more responsive to your other applications and your interactions.

VIRGINIA TECH

High Performance Computing

✓ ARC hosts several systems designed for high-performance and/or high-throughput computing (HPC/HTC)

Cluster	Description	Since
CUI	Dense GPU + some CPU for projects with controlled data/software	c. 2021
Tinkercliffs	HPC/HTC Flagship CPU HPE Dense GPU nodes (A100) DGX Dense GPU nodes (A100)	c. 2020 c. 2021 c. 2022
Infer (nearing end of life)	Accelerating inference and ML workloads (T4 GPU) Added P100 GPUs from Newriver Added V100 GPUs from Cascades	c. 2021 c. 2016 (EOL) c. 2018 (EOL)
OWL (coming soon)	Water-cooled latest generation AMD CPU high mem-per-core DDR5	c. 2024
Falcon (later in 2024)	GPU node expansion L40S GPUs (20 nodes x4 GPUs) A30 GPUs (32 nodes x4 GPUs)	c. 2024

TinkerCliffs - Flagship CPU Cluster

316 Nodes w/ 128 cores(AMD EPYC Rome)
16 Nodes w/ 96 cores (Intel Cascade Lake-AP)
41,984 CPU cores

tc-hm[001-008]
largemem_q

tc[001-308]
dev_q, preemptable_q
tc[001-307]

normal_q
tc[001-302]

interactive_q
tc308

tc-intel[001-016]

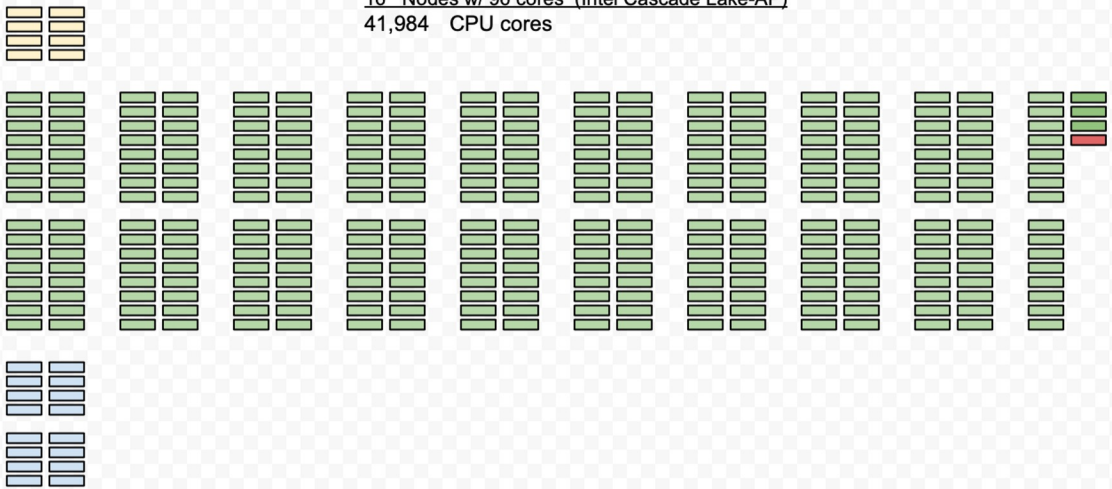

w/ dense GPU

ai[001-04]
a100_normal_q

4 Nodes w/ 128 cores (AMD Epyc Rome 7742)
+ 8 NVIDIA A100-80GB GPUs (6912 CUDA)
512 CPU cores
32 GPU accelerators
221,184 CUDA cores

ai[001-04]
a100_normal_q

10 Nodes w/ 128 cores (AMD Epyc Rome 7742)
+ 8 NVIDIA A100-80GB GPUs (6912 CUDA)
1280 CPU cores
80 GPU accelerators
552,960 CUDA cores

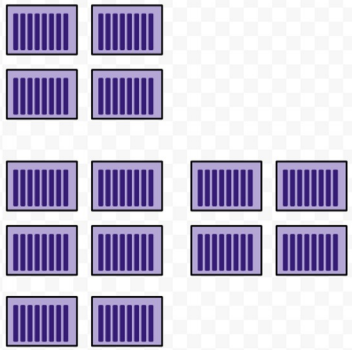

Soon: 2022 expansion

VIRGINIA TECH

Infer - Accelerating ML/DL and Inference

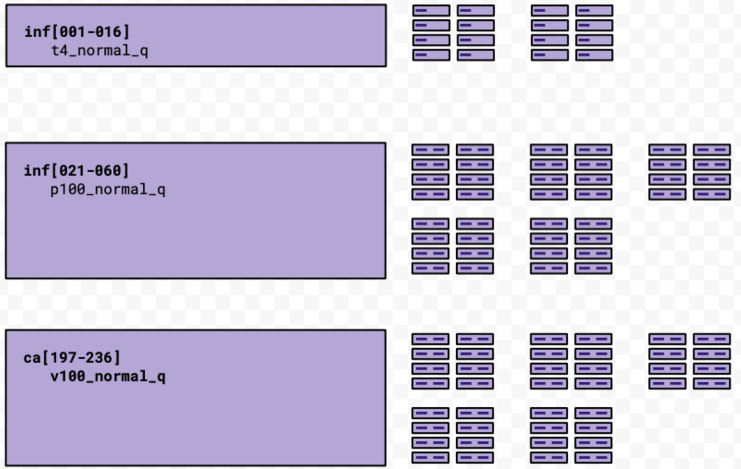

16 Nodes w/ 32 cores (Intel Skylake) + 1 NVIDIA T4 GPU (2560 CUDA + 320 tensor cores)
40 Nodes w/ 28 cores (Intel Broadwell) + 2 NVIDIA P100 GPUs (3580 CUDA cores)
40 Nodes w/ 24 cores (Intel Skylake) + 2 NVIDIA V100 GPUs (5,120 CUDA cores, 640 tensor cores)
2,592 CPU cores
176 GPU accelerators
593,760 CUDA cores
56,320 Tensor cores

CUI (Protected Data) System

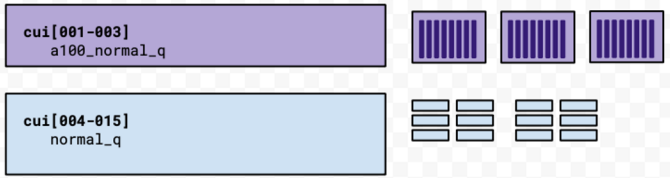

3 Nodes w/ 128 cores (AMD Epyc Rome 7742)
+ 8 NVIDIA A100-80GB GPUs (6912 CUDA)
12 Nodes w/ 64 cores + 512GB memory
1152 CPU cores
24 GPU accelerators
165,888 CUDA cores

VIRGINIA TECH

OWL - water cooled CPU with favorable memory architectures

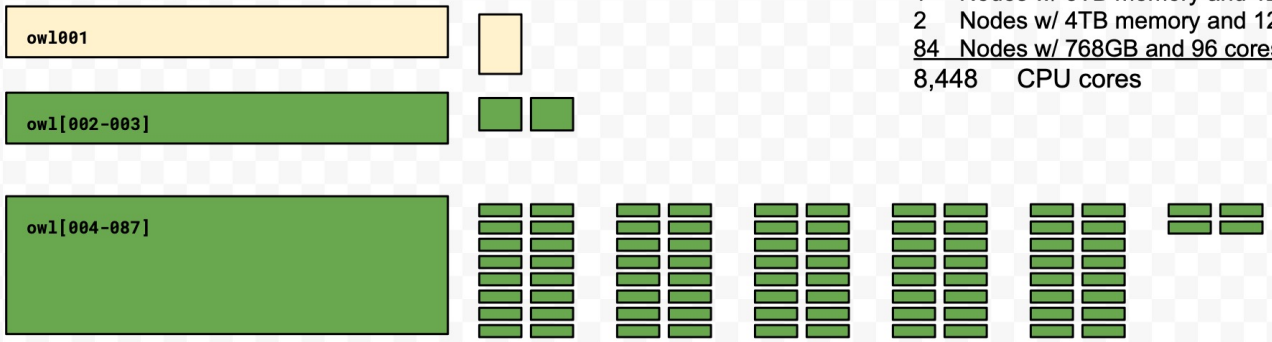

1 Nodes w/ 8TB memory and 128 cores (AMD Milan 7763)
2 Nodes w/ 4TB memory and 128 cores (AMD Milan 7763)
84 Nodes w/ 768GB and 96 cores (AMD Genoa 9454)
8,448 CPU cores

Falcon - mid-range GPU (Infer replacement)

20 nodes with 4x L40S GPU each (1.7x A100 AI Training Performance)

32 nodes with 4x A30 GPU each (5.2TF Peak FP64 each, 24GB)

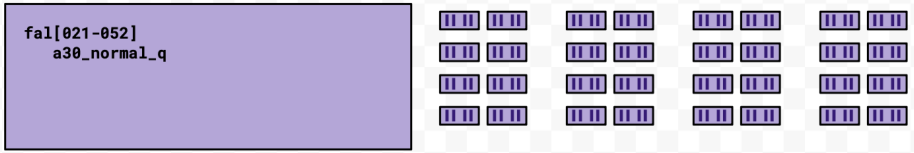

20 Nodes w/ 64 cores (Intel 8462Y+) + 4 NVIDIA L40S GPUs
32 Nodes w/ 64 cores (Intel 8462Y+) + 4 NVIDIA A30 GPUs
3,328 CPU cores
208 GPU accelerators

VIRGINIA TECH

Storage and Networks

Data storage systems:

HOME	personal files, low capacity, universal
PROJECTS	group shared storage, individual projects, large scale, universal
GLOBALSCRATCH	short term, staging jobs, 90-day aging
ARCHIVE	tape storage for data archival
LOCALSCRATCH	fastest I/O for jobs, wiped when job ends

Networks:

Campus Backbone & Datacenter network

100Gbps Infiniband interconnect – low latency.

Also 1, 10, 40, or 100Gbps Ethernet.

VPN needed for off-campus access.

https://www.docs.arc.vt.edu/resources/storage.html

VIRGINIA TECH

Systems

Aggregated computational resources:

- ✓ 500+ Compute nodes
- ✓ 50,000+ CPU cores
- ✓ 300+ GPU accelerators
- ✓ 10+ PiB data storage
- ✓ 900,000+ CUDA cores

+ high speed Ethernet and low-latency Infiniband interconnecting networks

+ large scale and high-performance storage systems

VIRGINIA TECH

Systems

Usage facts and figures:

2022-08-01 through 2023-07-31

- ✓ 1,138,442 Jobs submitted
- ✓ 292,539,937 CPU-hours allocated
- ✓ 1,196,697 GPU-hours allocated
- ✓ 1,194 Active users

Visualization

Visualization

- Desktop Visualization
- *HyperCube in the Visionarium Lab*
- User support and consulting
- Research collaboration
- Trainings and classes
- Tours and field trips

VIRGINIA TECH

https://arc.vt.edu/services/visualization.html

Industry standard usage model

- Linux compute clusters
  - ”headless” CentOS (moving to Rocky Linux) nodes
  - SLURM scheduler
  - Infiniband interconnect networks
  - EasyBuild software installation from source; modules approach to customizing computing environment per job type.
- Connect to “login node” using SSH client, upload/download files, command-line interface
- Running jobs
  - Compose job script, submit to scheduler (command line), job runs in batch mode on compute nodes
  - Interactive from command line.
  - OnDemand interactive session connections to clusters.

This remains a very productive model and the dominant mode of usage for many but can be a barrier to entry for others. ARC Helpdesk answers support tickets and hosts daily office hours to help.

High Level Operational Environment

Approach 1: Running jobs via batch processing

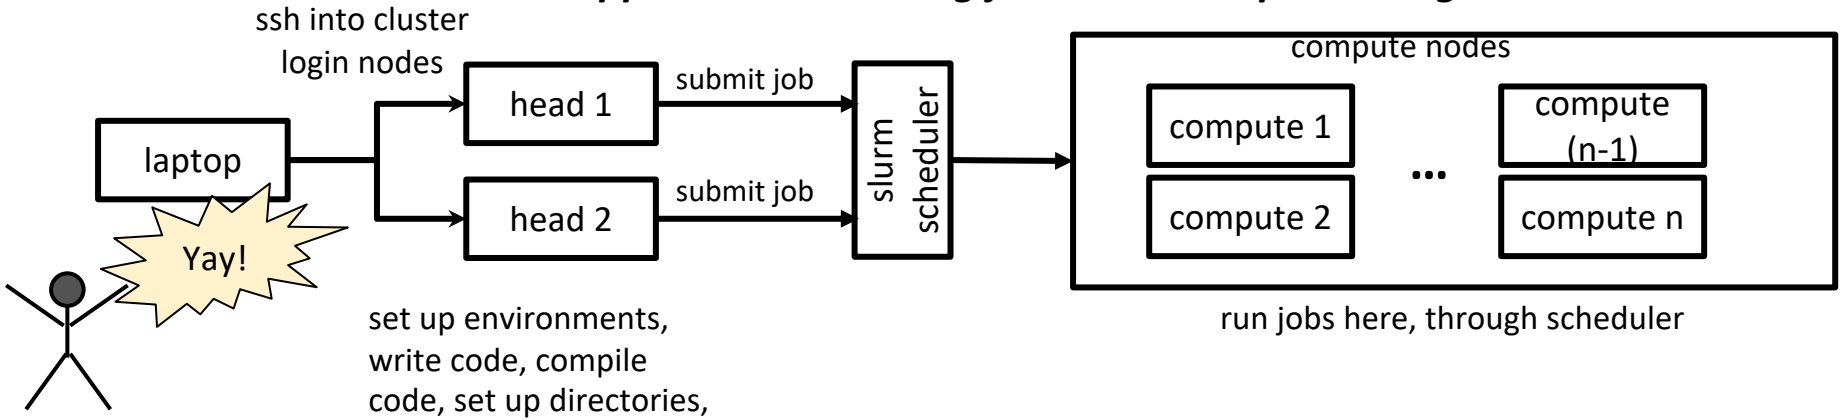

set up environments, write code, compile code, set up directories, write configuration files, create scripts, move files, submit jobs.

Do not run jobs on the head nodes, puhleeeze.

run jobs here, through scheduler

High Level Operational Environment

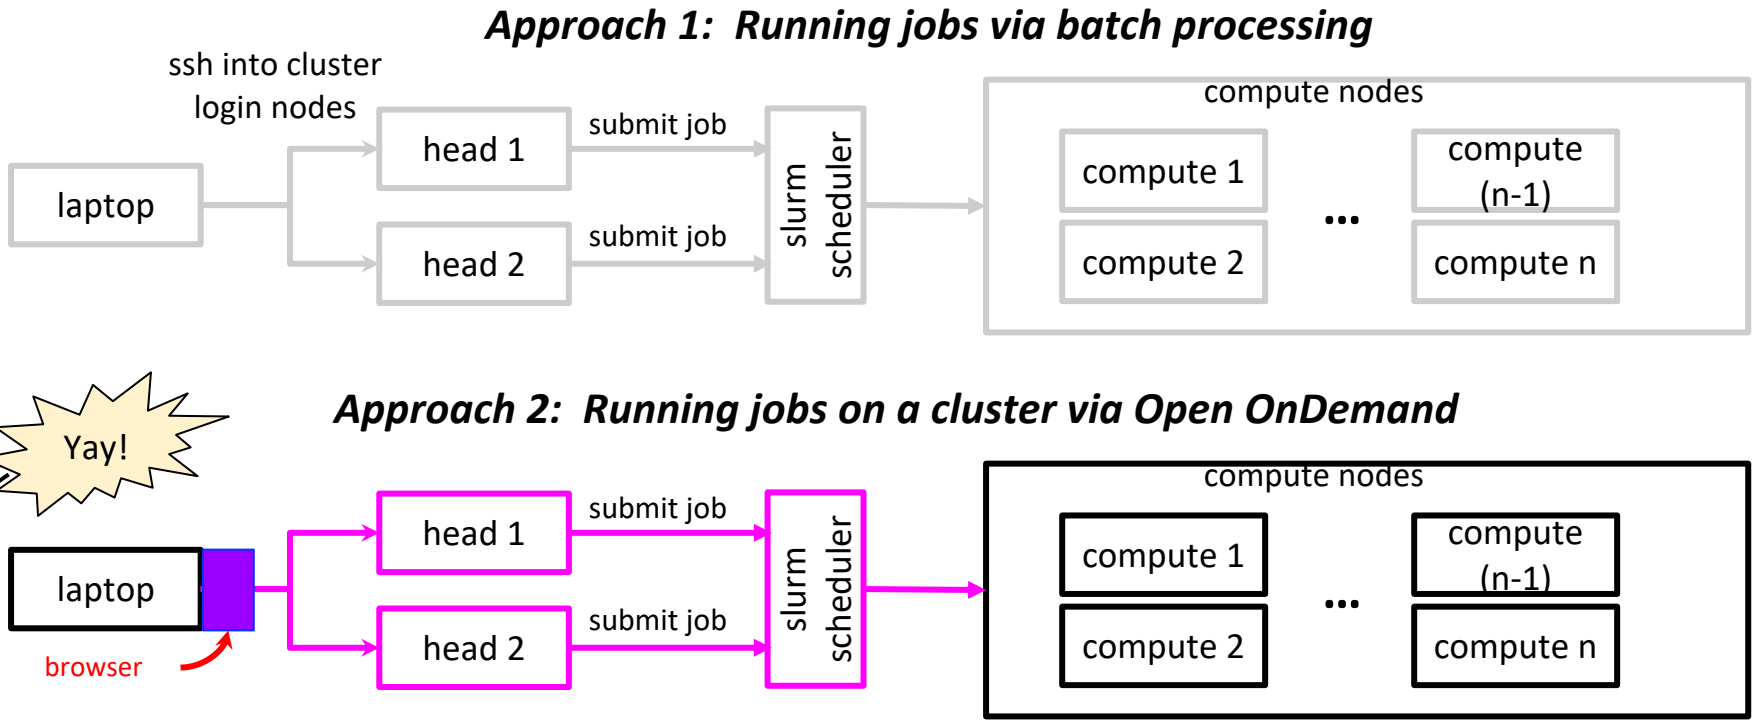

use your browser on your laptop to obtain compute nodes, and run jobs on them.

High Level Operational Environment

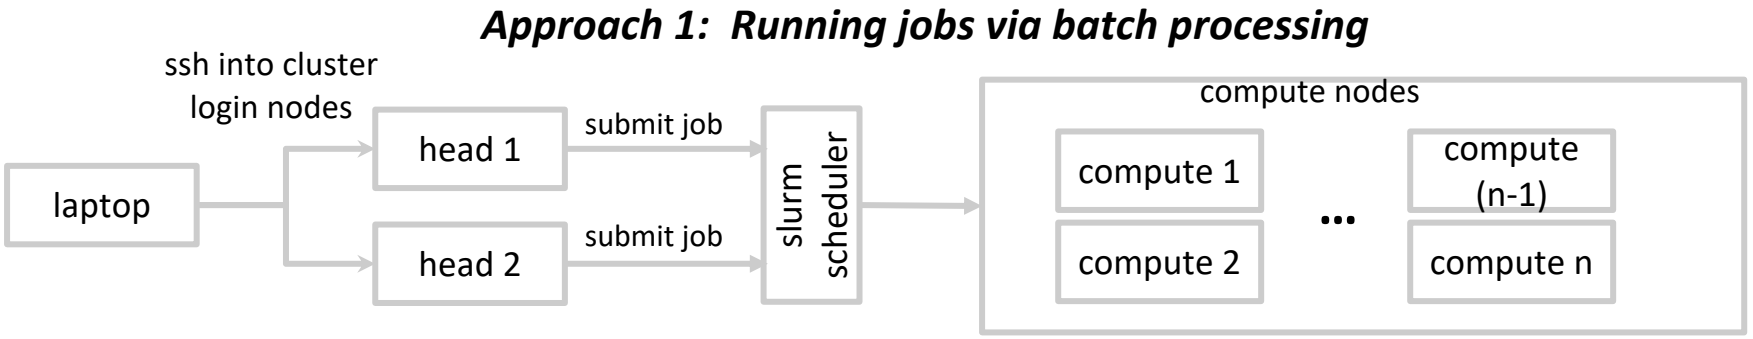

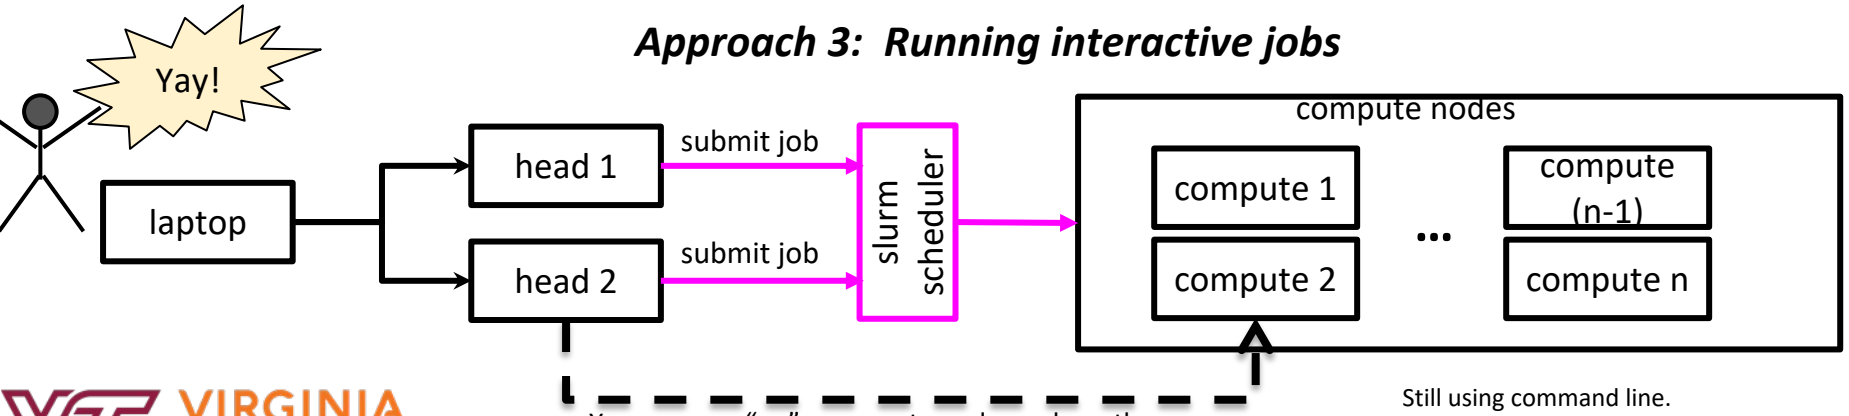

Removing barriers

Open OnDemand Web Interface

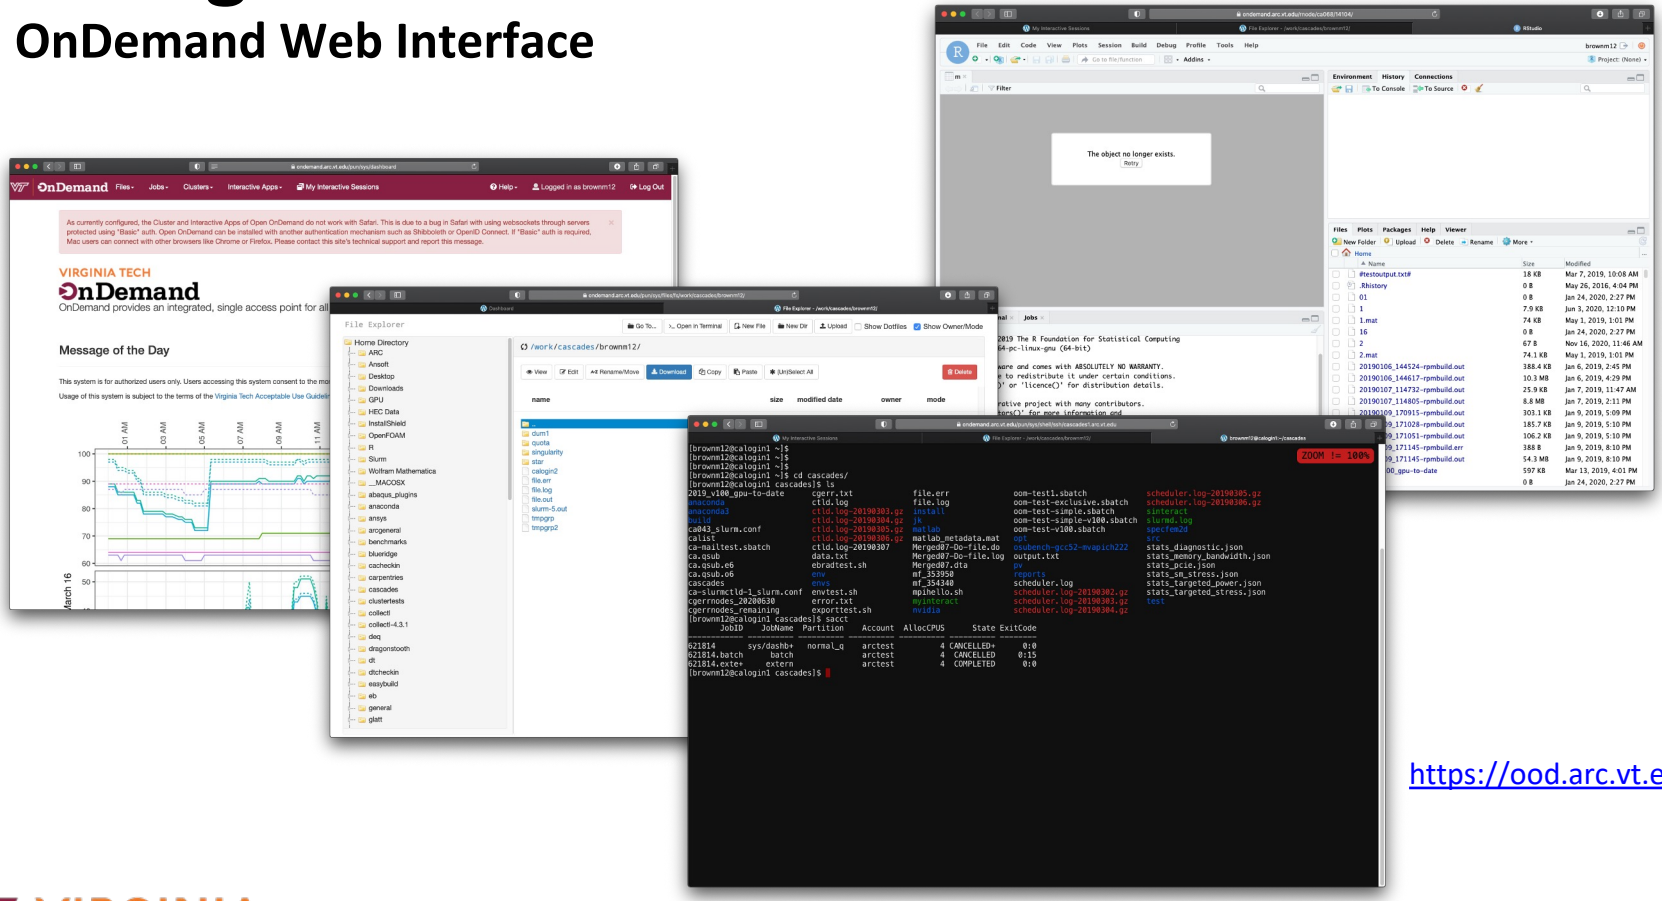

https://ood.arc.vt.edu

Removing barriers to entry

- Vast majority of ARC system usage is conducted at no direct cost to the researchers.
- Welcome all experience levels and fields of research.
- Provide state-of-the-art hardware and delivery models.
  - GPU accelerators for AI/ML/DL.
  - Many-core CPUs.
  - Support containerized software.
- Provide simplified interfaces wherever possible: Open OnDemand.

Resources and Services

Consultation / Collaboration / Helpdesk

Support, Consultation and Collaboration

ARC Documentation Website: https://docs.arc.vt.edu

ARC Helpdesk:

https://4help.vt.edu/sp?id=sc_cat_item_request&sys_id=4c7b8c4e0f712280d3254b9ce1050e3c

ARC Helpdesk GRAs work as a team to handle most incoming questions/problems.

"How do I setup SSH keys for authentication?" *"What can I do to get my job to launch faster?"* *"Why did my job stop?"*

"Is MATLAB available on Infer?" *"How can I share my files with my collaborator?"*

Office hours daily: https://arc.vt.edu/office-hours

GRAs escalate issues to ARC Computational Scientists as needed and meet bi-weekly as a group for collaborative discussions.

VIRGINIA TECH

Consulting and Collaboration

ARC Computational Scientists

- Have broad exposure to research applications and computational tools
- Provide research domain expertise
- Offer classes, short courses, and workshops
- Design workflows and assist with optimization of codes
- Build, install, and manage software on ARC systems
- Are the local experts on system design, software, and functionality
- Participate in research projects (co-author publications, co-PI on sponsored projects)
- Build research partnerships with centers, labs, projects, initiatives
- Want to engage very early in the proposal process to provide resources

Hardware and Systems Engineering

Engineering Activities

- Architect, install and maintain research system network, storage, compute resources and workload management
- Implement and maintain system security practices
- Respond to alerts from monitoring/logging
- Track and maintain physical assets and facilities
- Collaborate with internal stakeholders, such as the National Security Institute by managing a CUI system, and external entities, such as VA DEQ
- Research new/emerging technologies to integrate into ARC's systems
- Maintain user accounts
- Operate user facing systems such as ColdFront and Open OnDemand
- Implement and maintain supporting infrastructure systems and services

Engineering Team

- Jessie Bowman
- Miles Gentry
- Jeremy Johnson
- Nathan Liles
- William Marmagas
- Doug McMaster
- Ben Sandbrook

https://arc.vt.edu/about/our-team.html

Cost Center and Investment Computing

Generous "**Free Tier**" (VT subsidized) which satisfies needs of majority of projects using ARC

- Tinkercliffs: 800,000 units monthly (core hours) per user
- Projects Storage: 25TB storage per PI

Cost Center available on Tinkercliffs and newer systems for expanded usage + priority, pay for usage

Investment Computing to purchase dedicated access to resources

https://www.docs.arc.vt.edu/pi_info/costcenter.html

Resources and Services

Teaching / Workshops / Instruction

ARC Outreach and Educational work

- Guest lecture in regular courses.
- Occasionally instructor of record for sections.
- Give presentations (like this one) at departmental meetings.
- Conduct short courses and workshops via TLOS (Technology-enhanced Learning and Online Strategies) PDN.
- Organize focused discussions for research labs.
- Participate in Software Carpentries curriculum and instruction.
- Participate in regional and national communities of practice.
    - SuperComputing
    - PEARC
    - MARIA
    - WHPC
    - ACM …

ARC Course Offerings

Introduction to Advanced Research Computing

- This presentation (Zoom)

This workshop provides an informational overview of Virginia Tech's Advanced Research Computing (ARC) which provides centralized computational resources including high performance computing (HPC) systems to enable research at VT.

The content is intended for VT faculty, researchers, and students who are interested in hearing why ARC exists, what ARC has to offer in terms of computational resources and services, and then provides information about getting started with ARC.

- ARC mission and goals
- Detailed description of resources and services hosted by ARC
- Getting started: steps to set up your account and allocations
- Where to go for help and consultations

https://profdev.tlos.vt.edu/?query=arc

VIRGINIA TECH

ARC Course Offerings

Connect to ARC Systems and Run your First Jobs

- Monday, 03 June 2024 (Zoom)

This workshop is geared towards VT faculty, researchers, and students who are new to ARC. The aim is to provide orientation to the user-facing components of ARC systems and to demonstrate common connection and usage patterns.

This includes an overview and demonstration of ARC's web-based portal (Open OnDemand), and also how to connect with command-line oriented tools. Attendees with ARC accounts can follow along in a walkthrough of the most useful scheduler-interaction commands and an overview of building and submitting a sample workload, to the scheduler in the form of a batch job.

- Connect via Open OnDemand
- Connect via SSH
- Cluster and Scheduler Orientation
- Run Demo Jobs

https://profdev.tlos.vt.edu/?query=arc

VIRGINIA TECH

ARC Course Offerings

Running code/software on ARC Systems in Different Ways

- Tuesday, 04 June 2024 (Zoom)

ARC systems run software which spans the full spectrum of modern research computing. Many fields have evolved their software in various ways, but most often within the support models of research computing centers like ARC. This workshop addresses several of the most common software delivery models and how they can be accessed and used on ARC systems.

The demonstrations will be predominantly via the linux shell command line interface and will cover our "software modules" system, python environments via Anaconda, and also the main components needed for building software from source codes, particularly MPI software.

- Relevance of the environment and using interactive shell jobs
- Search for, load, and manage modules
- Python with Anaconda Environments
- Building software from source code

https://profdev.tlos.vt.edu/?query=arc

ARC Course Offerings

Launching Jobs in Parallel on ARC Clusters

- Tuesday, 04 June 2024 (Zoom)

The course delves into the details of parallel job execution, enabling participants to efficiently distribute computational workloads and maximize the utilization of ARC clusters.

This course is ideal for researchers, scientists, engineers, and computing professionals who want to leverage the capabilities of HPC clusters to accelerate their computations. Participants should have a basic understanding of programming concepts and a familiarity with Linux environments.

- Parallel Programming Models Background: Distinguish between parallel programming models
- Parallel Job Launching: Discover strategies for launching parallel jobs, considering factors like workload distribution and communication patterns. Practice launching parallel tasks using MPI launchers, Slurm launchers, and standard GNU parallel launchers
- Delve into the concept of hybrid parallel computing, combining multiple parallel programming models for enhanced performance.

https://profdev.tlos.vt.edu/?query=arc

VIRGINIA TECH

ARC Course Offerings

Monitoring Resource Utilization and Job Efficiency

- Wednesday, 05 June 2024 (Zoom)

Learn how to monitor and analyze the performance and efficiency of the computational jobs you run on ARC systems. Understanding the inter-relations of CPU utilization, memory utilization, I/O demand, and GPU utilization helps assess and organize efficient computational structures.

A variety of tools can be employed to including the command line tools: "seff", "jobload", "htop", "gpumon", "sacct", and more.

- Standard metrics which reflect the efficiency and performance of a workload.
- How Slurm job resource requests translate into CPU, memory, and GPU allocations for a job.
- Using tools to assess the performance of a workload while it is running.
- Familiarity with tools which can assess for completed jobs how efficiently the allocated resources were used.

https://profdev.tlos.vt.edu/?query=arc

Getting Started

Accounts / Accounting / Planning / Lifecycle

Getting Started

https://www.docs.arc.vt.edu/get_started.html

Needs Assessment

- Compute
- Storage
- Software
- Collaboration
- Visualization
- Lifecycle and data retention

Get an account

https://arc.vt.edu/account

- Get account for log-in

Register a Project and Get Allocations

https://coldfront.arc.vt.edu

- Create a “project”, add people, grants/pubs
- Request allocation for Compute to run jobs
- Request allocation for Project storage if desired

Where to get help

Website (https://docs.arc.vt.edu)

- FAQs
- Video demos
- Detailed instructions
- Examples

https://github.com/AdvancedResearchComputing/examples

Helpdesk (https://arc.vt.edu/help); then click "Request this service" button.

Office Hours (https://arc.vt.edu/office-hours)

Ask for consultation

- Workflow design
- Optimization
- Sponsored Projects

VIRGINIA TECH

Getting Assistance

Websites / Helpdesk / Office Hours / Consultation

Thanks for watching and listening!

ARC Website: www.arc.vt.edu

ARC documentation: https://www.docs.arc.vt.edu/get_started.html

My contact info: Sarah Ghazanfari
sarahghazanfari@vt.edu

Course Feedback:

https://docs.google.com/document/d/1RwLkcQgc1eDmIUjICXI99WWY2DGhoZgvDbK3VDffP1c/edit?usp=drive_link

VIRGINIA TECH