August 2023

Advanced Research Computing (ARC) Overview

Matthew Brown
Computational Scientist
Advanced Research Computing

Sign-In

VIRGINIA TECH

Fall 2023 ARC workshop Series

Introductory/Orientation Workshops (75 minutes each):

- Advanced Research Computing (ARC) Overview
- Connect to ARC systems and run your first jobs
- Get your software/code to run on ARC

Series 1: 8/17 - 8/18
Series 2: Wednesday afternoons: 9/27, 10/4, 10/11

Special Topics (75 minutes each):

- Monitoring Resource Utilization and Job Efficiency
- Getting the Best Data Storage Performance on ARC Filesystems
- Launching in Parallel

Series 1: Biweekly, Tuesday mornings 9/5, 9/19, 10/3
Series 2: Weekly, Thursday afternoons 10/26, 11/2, 11/9

VIRGINIA TECH

Expectations for Today

- This is an informal workshop
- Mostly informational about ARC and research computing at VT
- I want to hear your questions
- Welcome to use chat to ask questions + some time at the end
- Feedback needed to help improve future workshops
    - One up / one down at the end

VIRGINIA TECH

Scenarios bringing people to ARC

Scaling out: “Our team just completed the first run of our analysis and found that it took four hours to run on a laptop. The results are perfect, but we have 8,500 more of these to run and need finish in a few months.”

Scaling up: “I have an 80GB data set that I need to process using a colleague’s program. I have done this with 3GB data sets in the past, but my computer crashes when I try to process the larger data set. I think I need more memory.”

Platform for novel technologies: “I want to try out using this neural network to see if it provide insights into my problem. But training it on my data is taking weeks.”

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.

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

The ARC Team

VP for Research Computing:	Dane Skow
Visualization:	Nicholas Polys, Director Ben Sandbrook
Network Research Manager:	Mark K. Gardner
Computational Science:	Matthew Brown, Ayat Mohammed *(+2 open hires in progress)*
Systems Engineering/Administration/Development:	Miles Gentry, Jeremy Johnson, Doug McMaster, William Marmagas, Jessie Bowman, Nathan Liles

Plus our student interns and Helpdesk GRAs!

VIRGINIA TECH

Research Examples

380+ active projects each year, thousands of registered user accounts.

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, … 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"

VIRGINIA TECH

Research Examples

380+ active projects each year, thousands of registered user accounts.

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, … 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 / Visualization
  - Consultation / Collaboration / Helpdesk
  - Teaching / Workshops / Instruction
- Getting started
  - Accounts / Accounting / Planning / Lifecycle
  - Walkthrough
- Getting assistance
  - Websites / Helpdesk / Office Hours / Consultation

VIRGINIA TECH

Resources and Services –

High Performance Computing / High Throughput Computing / Visualization

VIRGINIA TECH

High Performance Computing

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

CUI	Dense GPU + some CPU for projects with controlled data/software	c. 2021
Tinkercliffs	HPC/HTC, Flagship CPU, Cost Center Capable AI/ML Dense GPU nodes more Dense GPU nodes	c. 2020 c. 2021 c. 2022
Infer	Accelerating inference and AI workloads	c. 2020
OWL (coming soon)	Water-cooled, latest generation AMD CPU, high mem-per-core, DDR5	c. 2023
Coming next	GPU node expansion	c. 2024

VIRGINIA TECH

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

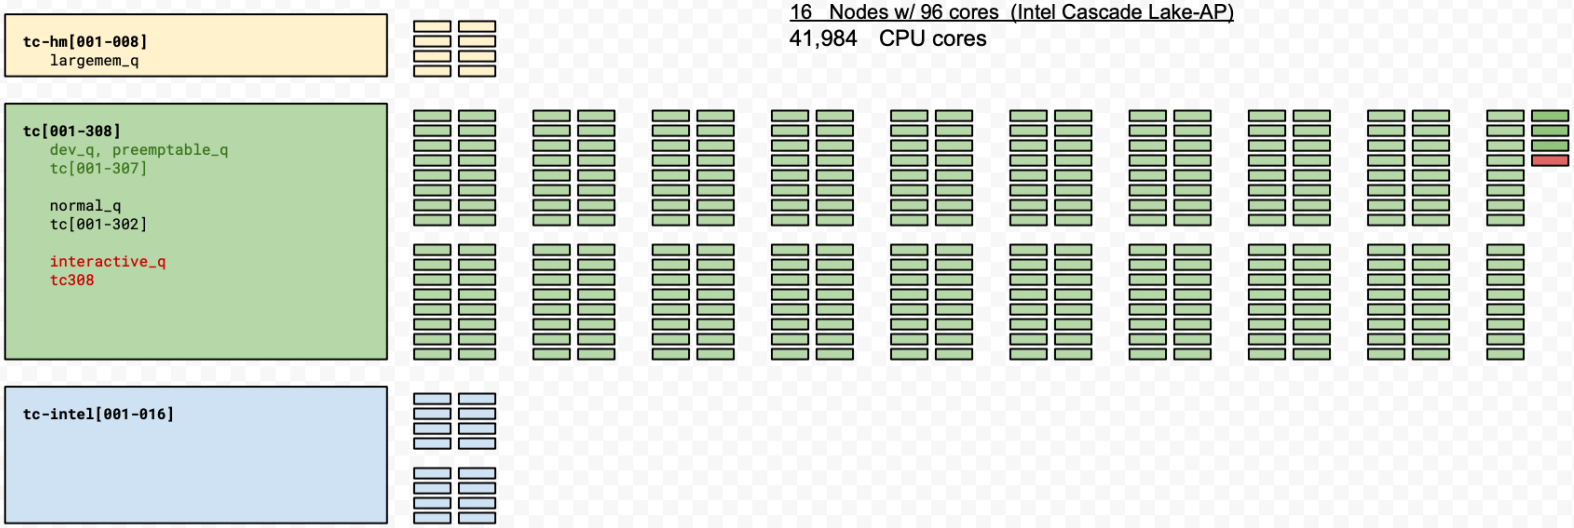

w/ dense GPU

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

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

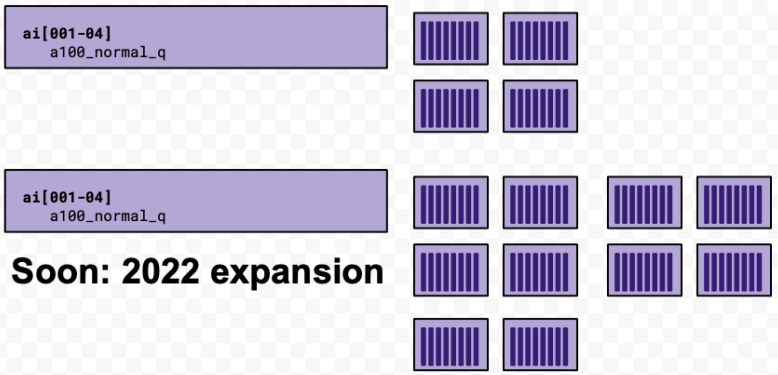

Soon: 2022 expansion

Infer - Accelerating ML/DL and Inference

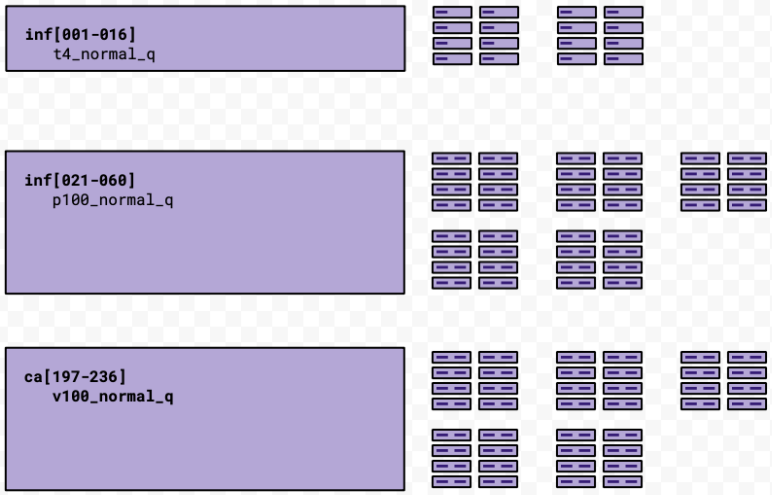

16 Nodes w/ 32 cores (Intel Skylake) + 1 NVIDIA T4 GPU (2560 CUDA + 320 tensor cores)
40 Nodes w/ 28 cores (Intel Broadwell) + 2 NVIDIA P100 GPUs (3580 CUDA cores)
40 Nodes w/ 24 cores (Intel Skylake) + 2 NVIDIA V100 GPUs (5,120 CUDA cores, 640 tensor cores)
2,592 CPU cores
176 GPU accelerators
593,760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

CUI (Protected Data) System

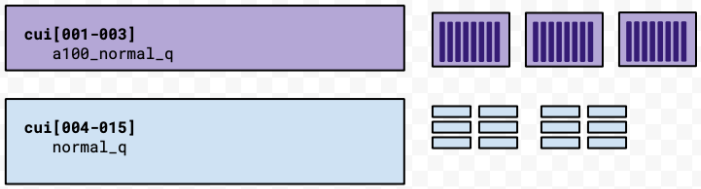

3 Nodes w/ 128 cores (AMD Epyc Rome 7742)
+ 8 NVIDIA A100-80GB GPUs (6912 CUDA)
12 Nodes w/ 64 cores + 512GB memory
1152 CPU cores
24 GPU accelerators
165,888 CUDA cores

VIRGINIA TECH

Storage and Networks

Data storage systems:

HOME	personal files, low capacity, universal
PROJECTS	group shared storage, 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:

450+	Compute nodes
50,000+	CPU cores
300+	GPU accelerators
10+	PiB data storage

+ 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

VIRGINIA TECH

Visualization

Visualization

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

VIRGINIA TECH

Industry standard usage model

- Linux compute clusters
  - ”headless” CentOS (moving to Rocky Linux) nodes
  - SLURM scheduler
  - Infiniband interconnect networks
  - EasyBuild software installation from source
- Connect to “login node” using SSH client, upload/download files, command-line interface
- Compose job script, submit to scheduler, job runs in batch mode on compute nodes

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.

Removing barriers

Open OnDemand Web Interface

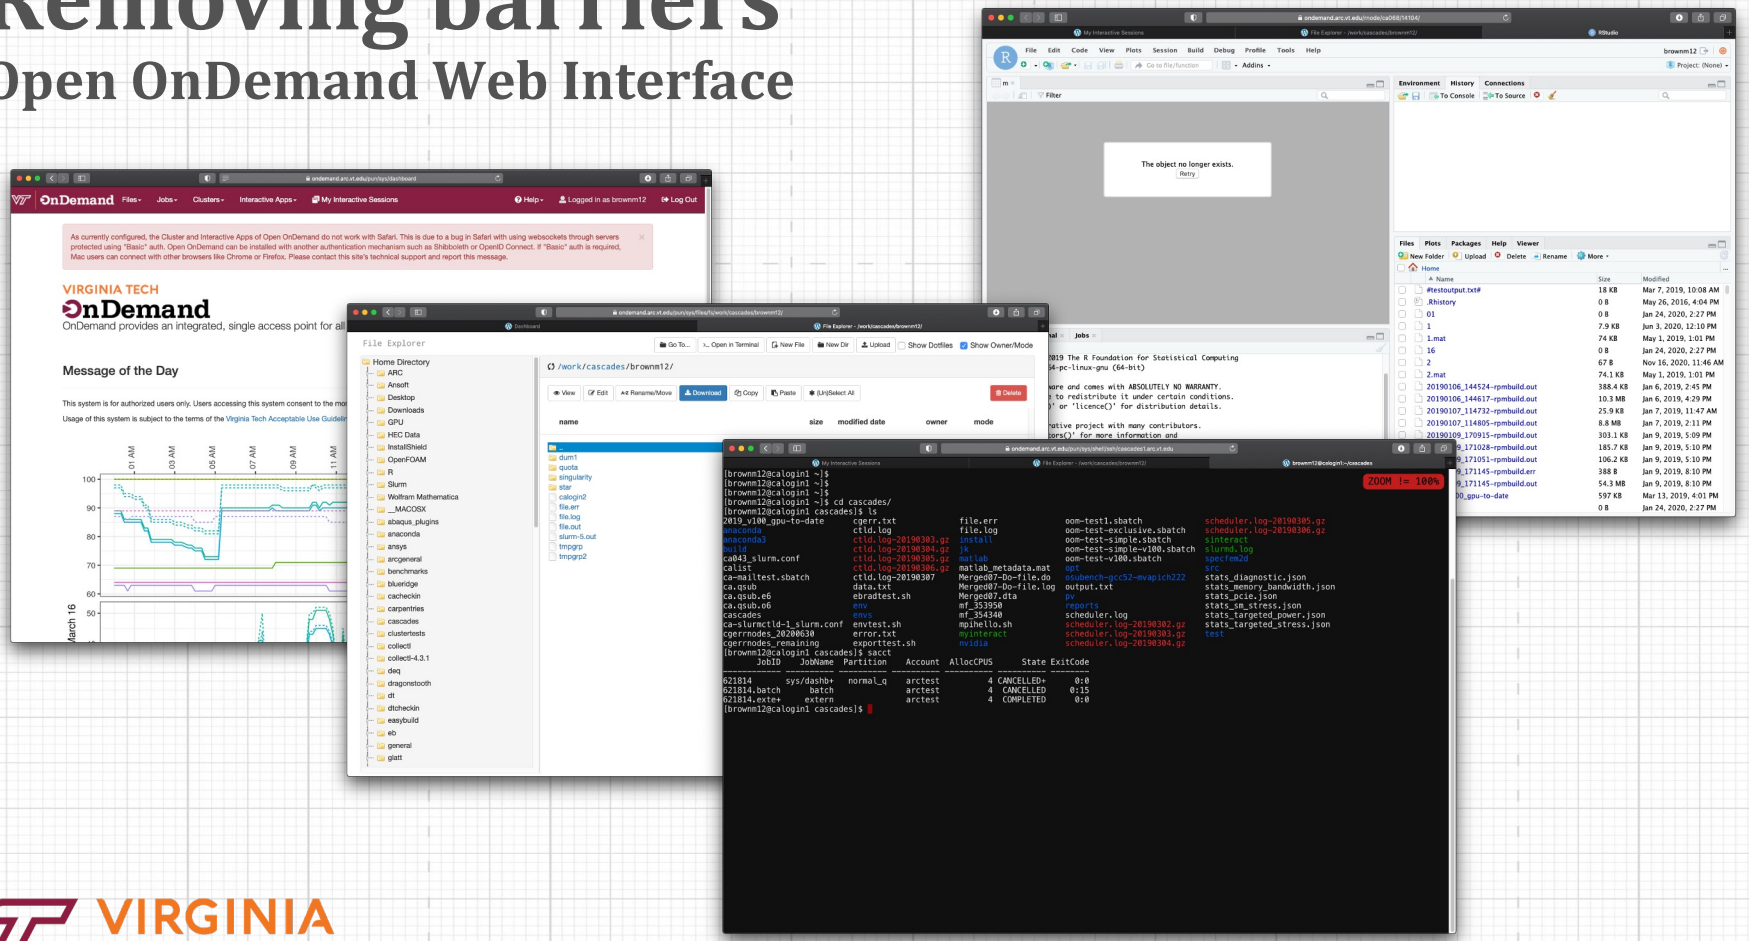

VIRGINIA TECH

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 of state-of-the-art hardware and delivery models
  - GPU accelerators for AI/ML/DL
  - Support containerized software
- Provide simplified interfaces wherever possible: Open OnDemand

VIRGINIA TECH

Resources and Services –

Consultation / Collaboration / Helpdesk

VIRGINIA TECH

Support, Consultation and Collaboration

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

ARC Helpdesk: https://arc.vt.edu/support

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 to 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

VIRGINIA TECH

Cost Center and Investment Computing

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

- Tinkercliffs: 600,000 units monthly (core hours)
- 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

VIRGINIA TECH

Huckleberry: Joint acquisition with ECE – Faculty Recruitment and Retention – Student Recruitment

- IBM "Minsky" compute node: higher performance per node, increase research productivity per unit time from a physical rack
- Fourteen nodes, IBM Power8 CPU + 4 NVIDIA P100 GPU per node
- Data movement limits performance for most applications — NVLINK accelerates data transfers within a node
- Deep learning applications have the largest performance advantages
- IBM Power AI software supports major open source deep learning frameworks optimized for POWER + NVIDIA architecture

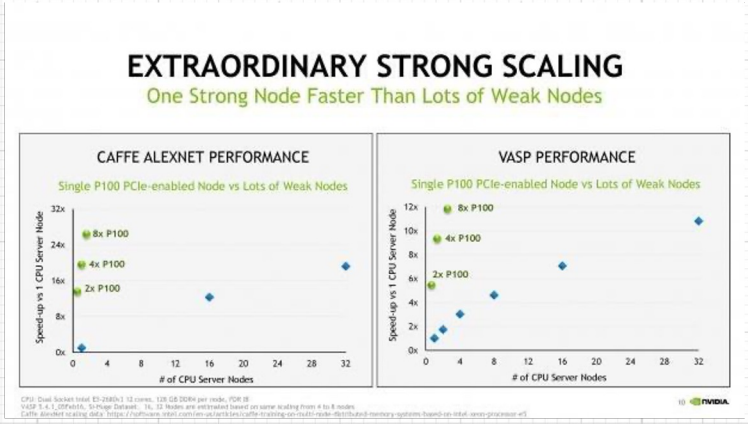

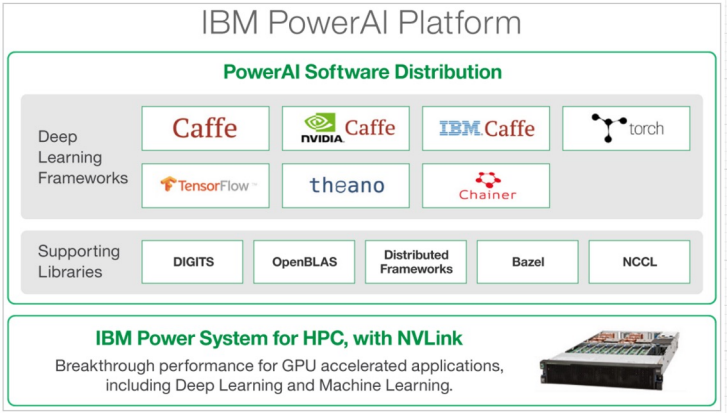

Image source: NVIDIA, IBM

VIRGINIA TECH

Controlled Unclassified Information: Joint acquisition with NSI/Hume – Expands capabilities for projects with controlled data

CUI system:

- Collaboration with VT-OESRC for controls
- High performance, scalable storage platform
- 24 powerful GPUs for AI, machine learning, and HPC
- 12 large CPU nodes

CUI is unclassified information requiring protection as identified in a law, regulation, or government-wide policy.

- The CUI Registry provides information on the specific categories and subcategories of information that the Executive branch protects. The CUI Registry also provides the newly approved Defense category. The CUI Registry can be found at: https://www.archives.gov/cui

CUI includes, but is not limited to:

- Privacy (including Health)
- Tax
- Law Enforcement
- Critical Infrastructure
- Controlled Technical Information
- Financial
- Intelligence
- Privilege
- Unclassified Nuclear
- Procurement and Acquisition

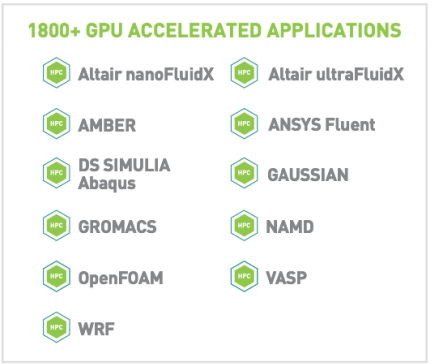

Image sources: NVIDIA, NIST

VIRGINIA TECH

Resources and Services –

Teaching / Workshops / Instruction

VIRGINIA TECH

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 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 …

VIRGINIA TECH

ARC Course Offerings

		Monday AM	PM	Tuesday AM	PM	Wednesday AM	PM	Thursday AM	PM	Friday AM	PM
Week 1	Aug 7 - Aug 11										
Week 2	Aug 14 - Aug 18							ARC1	ARC2	ARC3	
Week 3	Aug 21 - Aug 25	Classes Begin									
Week 4	Aug 28 - Sep 1										
Week 5	Sep 4 - Sep 8	Labor Day	Labor Day	SpecTop1							
Week 6	Sep 11 - Sep 15										
Week 7	Sep 18 - Sep 22			SpecTop2							
Week 8	Sep 25 - Sep 29						ARC1				
Week 9	Oct 2 - Oct 6			SpecTop3			ARC2			Fall Break	Fall Break
Week 10	Oct 9 - Oct 13						ARC3				
Week 11	Oct 16 - Oct 20										
Week 12	Oct 23 - Oct 27								SpecTop1		
Week 13	Oct 30 - Nov 3								SpecTop2		
Week 14	Nov 6 - Nov 10								SpecTop3		
Week 15	Nov 13 - Nov 17										
Week 16	Nov 20 - Nov 24					Thanksgiving Break		Thanksgiving Break		Thanksgiving Break	
Week 17	Nov 27 - Dec 1										
Week 18	Dec 4 - Dec 8							Reading Day	Reading Day	Exams	
Week 19	Dec 11 - Dec 15										
Week 20	Dec 18 - Dec 22										

VIRGINIA TECH

ARC Course Offerings

Advanced Research Computing (ARC) Overview

- Thursday, August 17th, 2023 9:30 - 10:45 (Zoom)
- Wednesday, September 27th, 2023 3:30 - 4:45 (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

VIRGINIA TECH

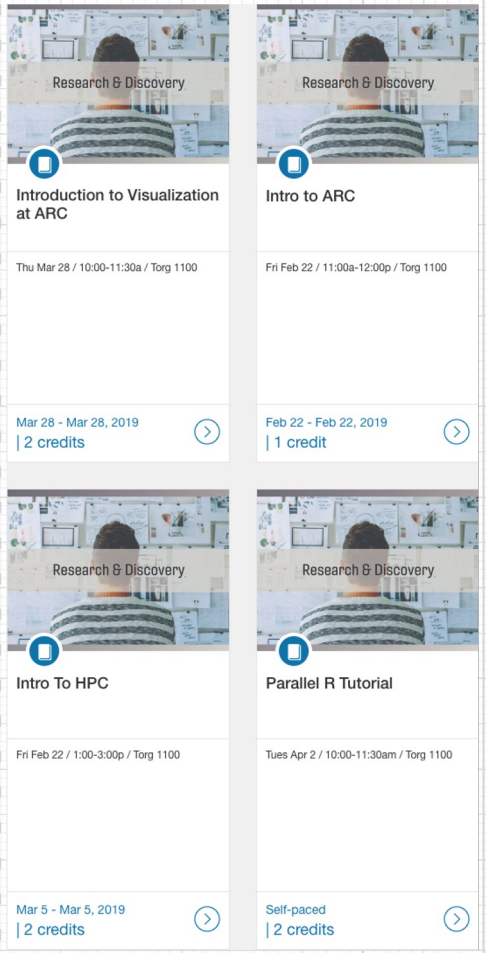

ARC Course Offerings

Connect to ARC systems and run your first jobs

- Thursday, August 17th, 2023 1:30 - 2:45 (Zoom)
- Wednesday, October 4th, 2023 3:30 - 4:45 (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

VIRGINIA TECH

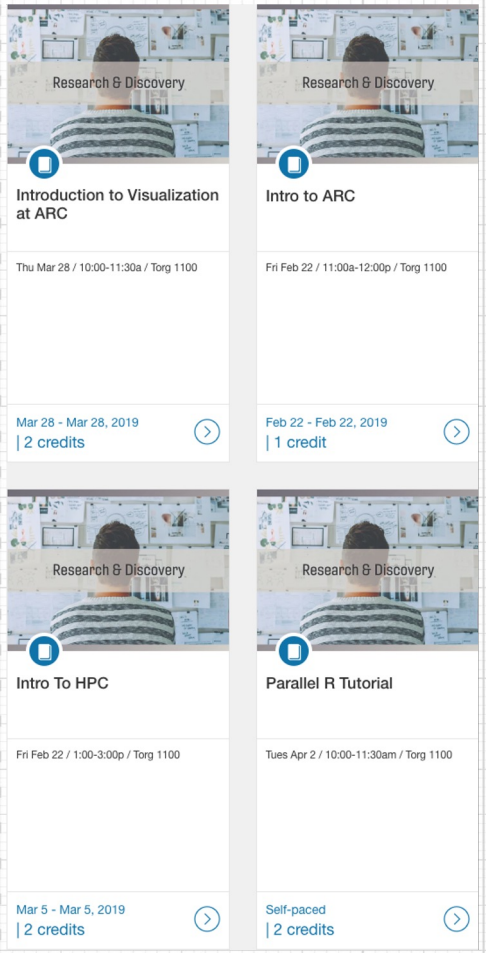

ARC Course Offerings

Get your software/code to run on ARC

- Friday, August 18th, 2023 9:30 - 10:45 (Zoom)
- Wednesday, October 11th, 2023 3:30 - 4:45 (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

VIRGINIA TECH

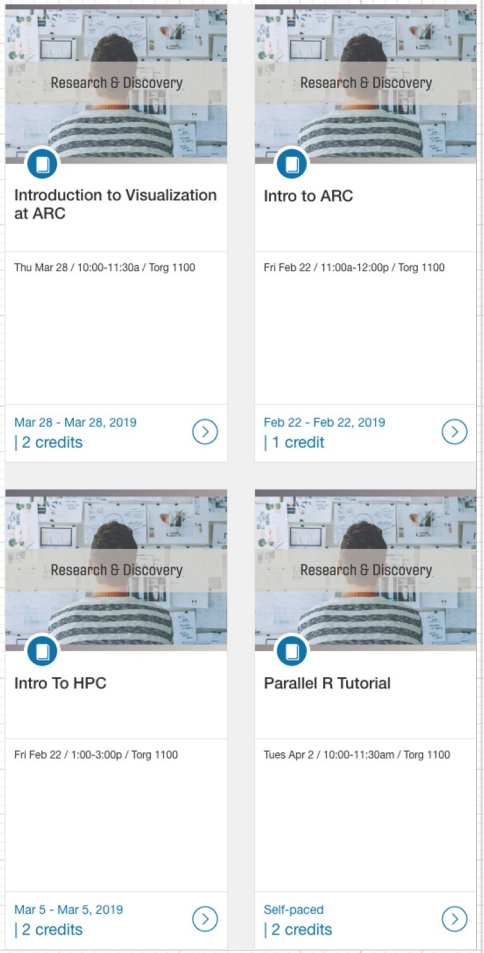

ARC Course Offerings

Monitoring Resource Utilization and Job Efficiency on ARC Systems

- Tuesday, September 5th, 2023 11:00 - 12:15 (Zoom)
- Thursday, October 26th, 2023 1:30 - 2:45 (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.

VIRGINIA TECH

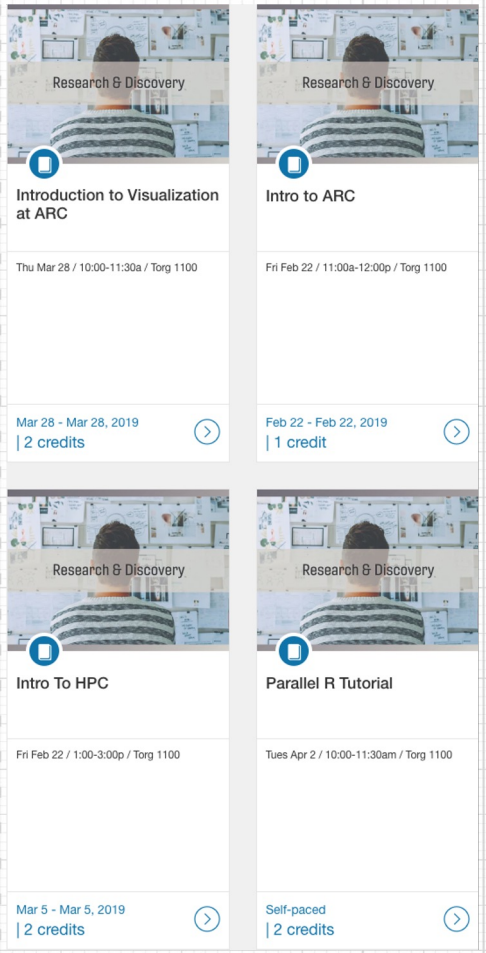

ARC Course Offerings

Getting the Best Data Storage Performance on ARC Filesystems

- Tuesday, September 19th, 2023 11:00 - 12:15 (Zoom)
- Thursday, November 2nd, 2023 1:30-2:45 (Zoom)

While CPU, memory, and GPU utilization often get the most attention, frequently data input-output (I/O) from storage becomes an unexpected bottleneck in scaling workloads up on clusters such as those hosted by ARC.

This workshop will explore the variety of ARC systems for file storage, discuss the strengths and design of each, and provide some best practices for using them. Be ready with your own use cases in mind as time will also be dedicated to discussing real-life example cases as brought by participants.

- Discuss pros and cons of network-based filesystems, local disks, and other storage targets.
- Policies and limits which govern usage of ARC storage systems
- Know how to select the right storage target for staging jobs vs. bulk storage vs. archival.
- Learn about resources available for managing data when working as a group.
- Practice using command line tools and system best practices for these and for moving data.

VIRGINIA TECH

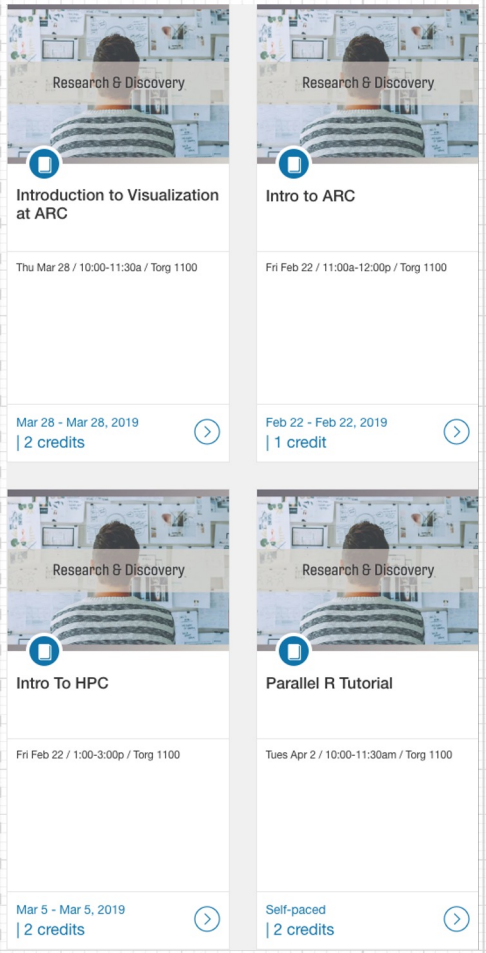

ARC Course Offerings

Launching in Parallel

- Tuesday, October 3rd, 2023 11:00 - 12:15 (Zoom)
- Thursday, November 9th, 2023 1:30 - 2:45 (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.

VIRGINIA TECH

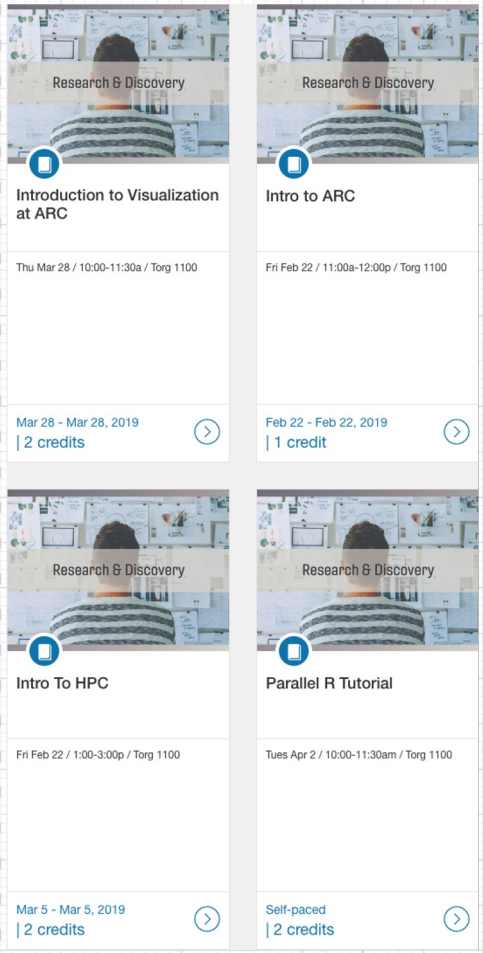

Getting Started –

Accounts / Accounting / Planning / Lifecycle

VIRGINIA TECH

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/exampl

Helpdesk (https://arc.vt.edu/help)

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

Ask for consultation

- Workflow design
- Optimization
- Sponsored Projects

VIRGINIA TECH

Getting Started –

Walkthrough

VIRGINIA TECH

Getting Assistance –

Websites / Helpdesk / Office Hours / Consultation

VIRGINIA TECH

Thanks for watching and listening!

ARC Website: www.arc.vt.edu

My contact info: Matthew Brown
brownm12@vt.edu

Course Feedback:
https://docs.google.com/document/d/1_Eaix0btZG3HKg6yEZiDUlDVSkJYPh6HxKKUGwirPZY/edit?usp=sharing

one up / one down

VIRGINIA TECH