April 2022

Advanced Research Computing (ARC) Overview

Matthew Brown, Computational Scientist
Advanced Research Computing
Information Technology
Virginia Tech

Sign-In

VIRGINIA TECH

Spring 2022 ARC workshop Series

April 12th or 13th	Advanced Research Computing (ARC) Overview	Mission and goals, resources and services, getting started, getting assistance
April 19th or 20th	Connect to ARC systems and run your first jobs	VPN, Windows Subsystem for Linux, Git/BASH, MobaXterm/PuTTY, OnDemand, ssh keys, screen/tmux
April 26th or 27th	Get your software to run on ARC	File management, finding things, monitoring utilization, understanding your environment, loading software
May 12th – 19th	Software Carpentries (VT Libraries)	Foundations of Unix, Git and Python. Programming with Python. R for Reproducible Research. The Unix Shell. Version Control with Git

VIRGINIA TECH

Expectations

- 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

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 a 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

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

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

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

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

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.

ARC’s Mission

In a nutshell our goals are to:

- Engage with researchers across all colleges, centers, and disciplines at Virginia Tech

… and to …

- Provide systems and support that advance research programs at VT

Who We Are

Associate VP for Research Computing:	Terry Herdman
Assist. Director, Development and Fiscal Admin:	Alana Romanella
Network Research Manager:	Mark Gardner
Director, Visualization:	Nicholas Polys
Computational Scientists:	Bob Settlage, Matt Brown, two open positions
Director of ARC Operations:	Kevin Shinpaugh
Systems Engineers/HPC System Administrators/ Software Engineers	Jeremy Johnson, Doug McMaster, William Marmagas, Jessie Bowman, Ben Sandbrook, Hunter Irving, Nathan Liles

Plus our student interns and Helpdesk GRAs!

http://careers.pageuppeople.com/968/cw/en-us/job/520041/computational-scientist

VIRGINIA TECH

Research Examples

380+ Active Projects. Faculty, adjunct faculty, postdocs, graduate students, undergraduate students.

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. Faculty, adjunct faculty, postdocs, graduate students, undergraduate students.

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 acting on an inclined, flexible, thin fin that is moving inside a fluid… applications … 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

High Performance Computing

ARC hosts a number of 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
Cascades	Heterogenous HPC: CPU/GPU/Large Memory	c. 2017, 2018
Dragonstooth	Moderate scale HTC	c. 2016

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

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]

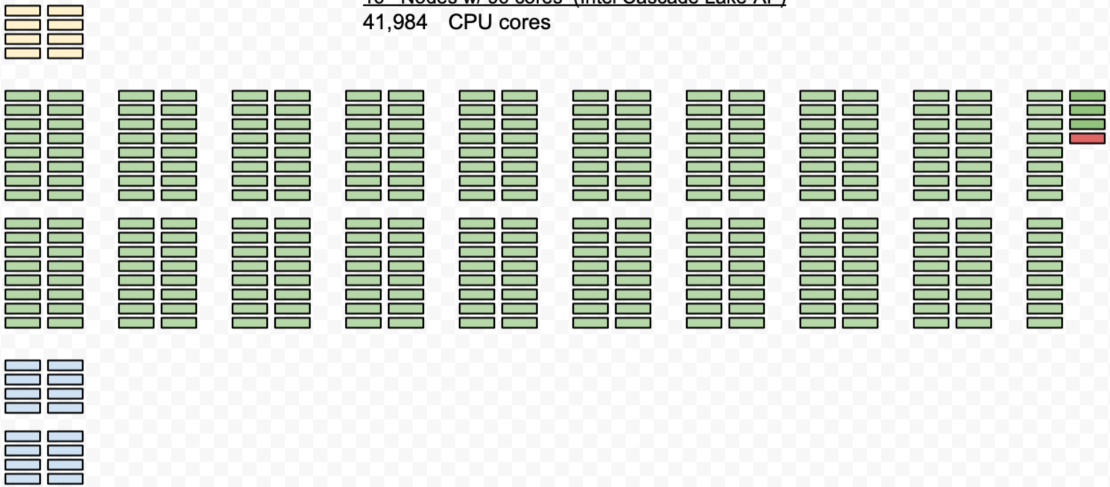

w/ dense GPU

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

Soon: 2022 expansion

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

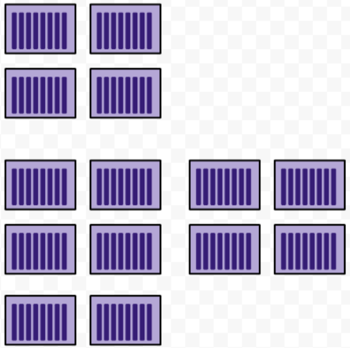

Infer - Accelerating ML/DL and Inference

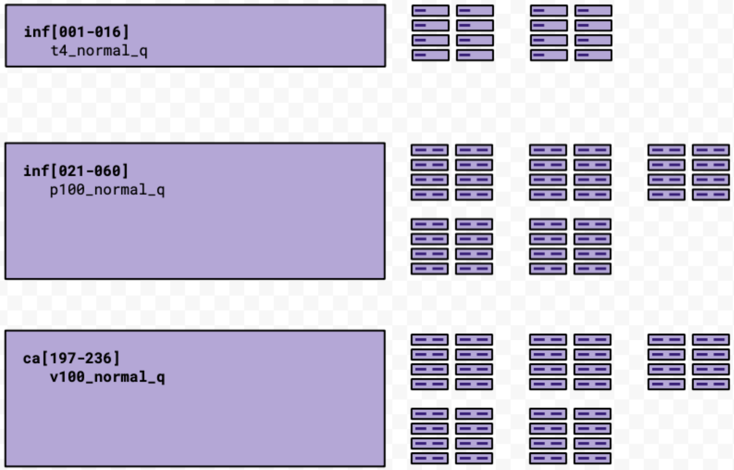

16 Nodes w/ 32 cores (Intel Skylake) + 1 NVIDIA T4 GPU (2560 CUDA + 320 tensor cores)
40 Nodes w/ 28 cores (Intel Broadwell) + 2 NVIDIA P100 GPUs (3580 CUDA cores)
40 Nodes w/ 24 cores (Intel Skylake) + 2 NVIDIA V100 GPUs (5,120 CUDA cores, 640 tensor cores)

2,592	CPU cores
176	GPU accelerators
593,760	CUDA cores
56,320	Tensor cores

CUI (Protected Data) System

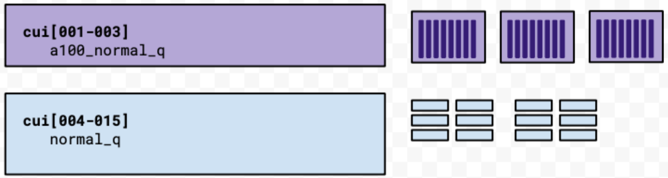

3 Nodes w/ 128 cores (AMD Epyc Rome 7742)
+ 8 NVIDIA A100-80GB GPUs (6912 CUDA)
12 Nodes w/ 64 cores + 512GB memory

1152	CPU cores
24	GPU accelerators
165,888	CUDA cores

Cascades - Heterogeneous HPC + GPU

```
ca[001-002]
   largemem_q
```

```
ca[003-006]
   k80_q
```

```
tc[007-196]
   dev_q, preemptable_q
   tc[007-196]

   normal_q
   tc[007-195]
```

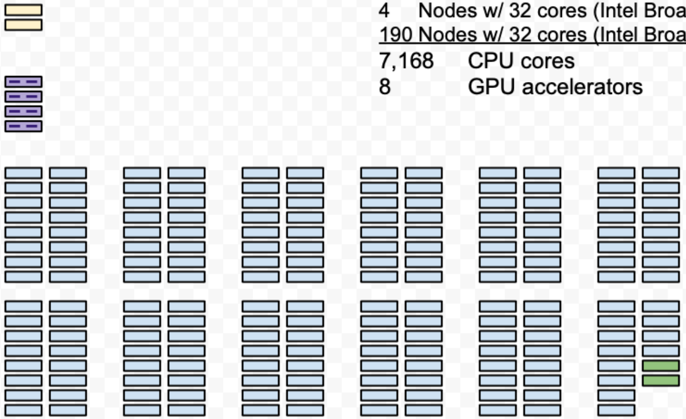

2 Nodes w/ 72 cores (Intel Broadwell)
4 Nodes w/ 32 cores (Intel Broadwell) + 2 NVIDIA K80 GPUs (4,992 CUDA cores)
<u>190 Nodes w/ 32 cores (Intel Broadwell)</u>

7,168 CPU cores
8 GPU accelerators

Dragonstooth - HTC

```
dt[003-048]
   normal_q
```

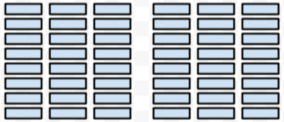

<u>48 Nodes w/ 24 cores (Intel Haswell)</u>

1,152 CPU cores

Scheduler permits very long jobs (30 days)
Scheduler permits large volumes of small jobs

VIRGINIA TECH

Systems

Aggregate resources:

701	Compute nodes
54,560	CPU cores
320	GPU accelerators
1,533,792	NVIDIA CUDA cores

+ high speed Ethernet and low-latency Infiniband interconnecting networks
+ large scale and high-performance parallel storage

VIRGINIA TECH

Storage and Networks

Data storage systems:

HOME	personal storage, low capacity, universal
PROJECTS	group shares, large scale
FASTSCRATCH	short term storage, staging jobs
ARCHIVE	tape storage for data archival
Local scratch	fastest I/O for jobs, wiped when job ends

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

Networks

Campus Backbone & Datacenter network

100Gbps Infiniband interconnect – low latency

Also 1, 10, 40, or 100Gbps Ethernet

VPN needed for off-campus access

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

Removing barriers to entry

- Vast majority of ARC system usage is consumed 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 (eg. GPUs for AI/ML/DL, containerization, cloud computing, etc.)
- Provide simplified interfaces wherever possible: Open OnDemand

  https://ood.arc.vt.edu

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

GRAs escalate issues to ARC Computational Scientists as needed.

VIRGINIA TECH

Consulting and Collaboration

ARC Computational Scientists

- Have exposure to many applications and software
- Provide research domain expertise
- Offer classes, short courses and workshops
- Design workflows and optimization of codes
- Build, install, and manage software on systems
- Are the local experts on system design, software, and functionality

Also…

- Provide Investment Computing Program and newly developed cost-center
- Can participate in research projects (co-author for publications, co-PI on grants)
- Build research partnerships with centers, labs, projects, initiaves
- 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
- Older ARC systems usage is unlimited (as available)

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

- $0.0023 per base core-hour on Tinkercliffs
- $2.1694 per TB/month storage

Investment Computing to purchase dedicated access to resources

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

VIRGINIA TECH

Huckleberry Joint 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

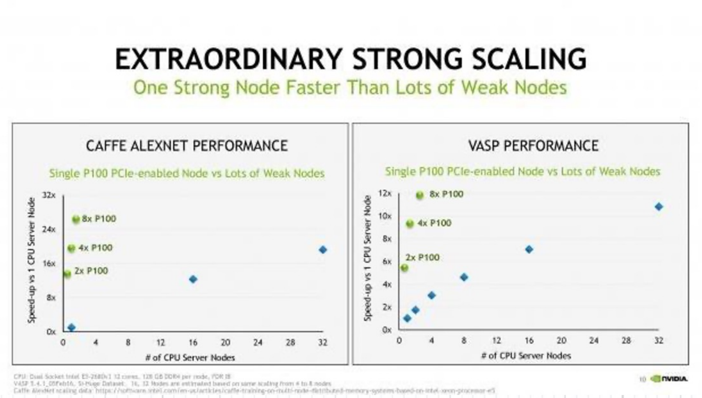

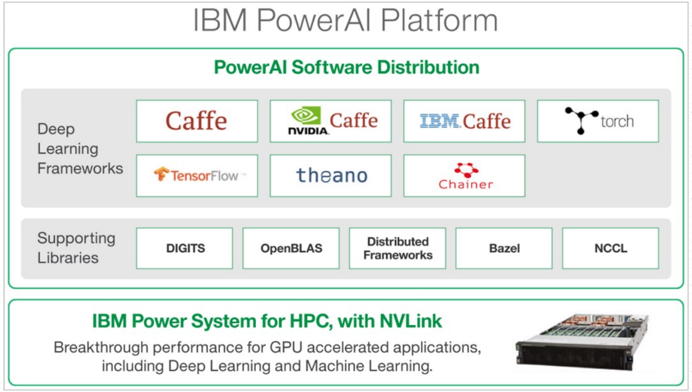

VIRGINIA TECH

Image source: NVIDIA, IBM

Controlled Unclassified Information Joint with NSI/Hume – Expands capabilities for projects with sensitive 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

EVERY DEEP LEARNING FRAMEWORK

Caffe2 · Microsoft Cognitive Toolkit · mxnet · PYTORCH · Apache Spark · TensorFlow · theano

1800+ GPU ACCELERATED APPLICATIONS

- Altair nanoFluidX
- Altair ultraFluidX
- AMBER
- ANSYS Fluent
- DS SIMULIA Abaqus
- GAUSSIAN
- GROMACS
- NAMD
- OpenFOAM
- VASP
- WRF

VIRGINIA TECH

Image sources: NVIDIA, NIST

Resources and Services –

Teaching / Workshops / Instruction

ARC Course Offerings

- ARC personnel offer or guest lecture in regular courses
- ARC personnel offer short courses and workshops
- ARC personnel participate in Software Carpentries workshops
- Examples:
  - Introduction to ARC systems
  - Connect to ARC systems and run your first jobs
  - Get your code or software to run on ARC
  - Deep learning with NVIDIA Digits
  - Python for scientific computing
  - Parallel R
  - Visual computing & Virtual Reality
  - The Unix Shell

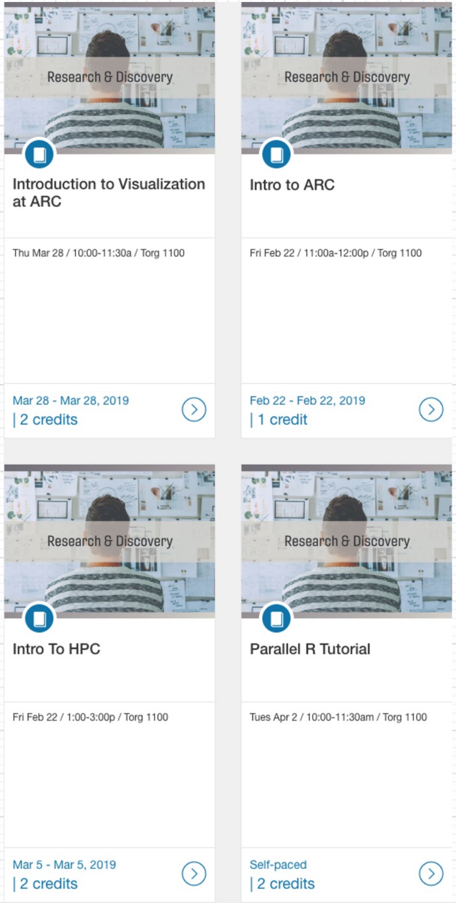

VIRGINIA TECH

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/examples

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

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

Request consultation

- Workflow design
- Optimization
- Projects

VIRGINIA TECH

Getting Started –

Walkthrough

Getting Assistance –

Websites / Helpdesk / Office Hours / Consultation

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