ABoVE Science Cloud on ADAPT

October 1st, 2021

Overview of the ABoVE Science Cloud on ADAPT

- ADAPT is the Advanced Data and Analytics PlaTform
- ABoVE Science Cloud (ASC) is a part of ADAPT
- Linux & Windows Virtual Machines (VMs)
- Some large staged datasets

Find out more here:
https://www.nccs.nasa.gov/systems/ADAPT

Linux VMs

- above101-103 (6 CPU, 60 GB RAM)
  - Deprecated Debian 8 nodes to be upgraded to CentOS 7.
- above104-110 (4 CPU, 48 GB RAM)
  - CentOS 7.9
  - Submit jobs via Slurm (above104 is the login node)
  - Connected to the ABoVE JupyterHub instance
    - https://www-proxy-dev.nccs.nasa.gov/jupyterhub-above/
- abovex101 (dedicated data transfer node)

How do I access the Linux VMs?

https://www.nccs.nasa.gov/nccs-users/instructional/adapt-instructional

Windows VMs

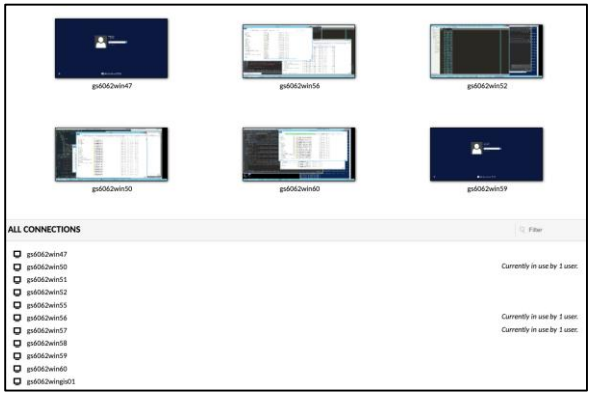

- 8 VMs
  - 14 CPUs
  - 104GB RAM
  - 1TB C:\ drive (used for temporary files only)
  - $HOME and $NOBACKUP shared between Linux and Windows
  - Connected to all other shared filesystems

How do I access the Windows VMs?

https://www.nccs.nasa.gov/nccs-users/instructional/adapt-instructional

How do I access the VMs?

// ACCESSING ADAPT VMS

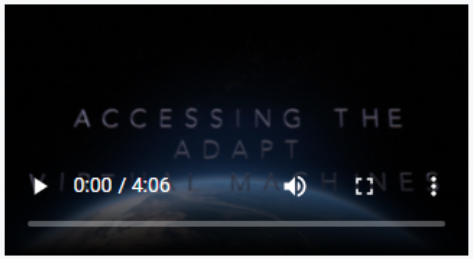

ACCESSING ADAPT INSTRUCTIONAL VIDEO

Learn how to access and log in to a virtual machine on the ADAPT system. Instructions include: 1) accessing ADAPT Linux VMs from external Linux systems; 2) accessing ADAPT Linux VMs from external Windows systems; and 3) accessing ADAPT Windows VMs directly through a web browser.

USER TYPES	ACCESSING WINDOWS VMs	ACCESSING LINUX VMs
WINDOWS USERS	Guacamole	MobaXterm application
MAC USERS	Guacamole	$ ssh adaptlogin.nccs.nasa.gov
LINUX USERS	Guacamole	$ ssh adaptlogin.nccs.nasa.gov

https://www.nccs.nasa.gov/nccs-users/instructional/adapt-instructional

Data Collections on ADAPT

- Available under the /css filesystem

- ABoVE: 80 TB
  - Oak Ridge National Laboratory Distributed Active Archive Center (ORNL DAAC) datasets, including over 130 archived as part of ABoVE.
  - National Snow and Ice Data Center (NSIDC) DAAC datasets, including LVIS products archived in support of ABoVE.
  - Alaska Satellite Facility (ASF) datasets related to L-band SAR.
- AMSR-2: 5 TB
- AVHRR/Polar: 40 TB on ADAPT and 10 GB on CSS
- CFHA: 250 TB
- CMIP5: 105 TB
- CREATE-IP: 79 TB
- CSDA-Spire: 30 TB
- DSCOVR: 72 TB (EPIC O3SO2AI)
- DSCOVR: 72 TB (L1B)
- DSCOVR: 72 TB (L2_CLOUD_03)
- FLDAS: 40 TB
- GeoMIP: 14 TB
- Geostationary (GOES): Ingest starting now, planning for 1 PB
- GEOS-5 Nature Runs (g5nr): 5 PB
- HIMAT Snow Reanalysis: 5 TB
- ICEBridge: 2 TB
- ICESat: 8 TB
- ICESat-2: 161 TB
- IMERG: 15 TB
- Landsat: 186 TB
- MAIAC: 107 TB
- MERRA: 86 TB
- MERRA2: 320 TB
- Selected MODIS data: 679 TB
- NEX GDDP: 11 TB
- NEX DCP30: 11 TB
- NGA: 4446 TB (~4.5 PB) *

https://www.nccs.nasa.gov/services/data-collections/css-collections

DigitalGlobe Imagery on ADAPT

- Signed NASA-NGA form is required
  - Only specified zones provisioned
- Catalog is now approximately 8.5 PB
  - 19.3 million images
- DEM data also available

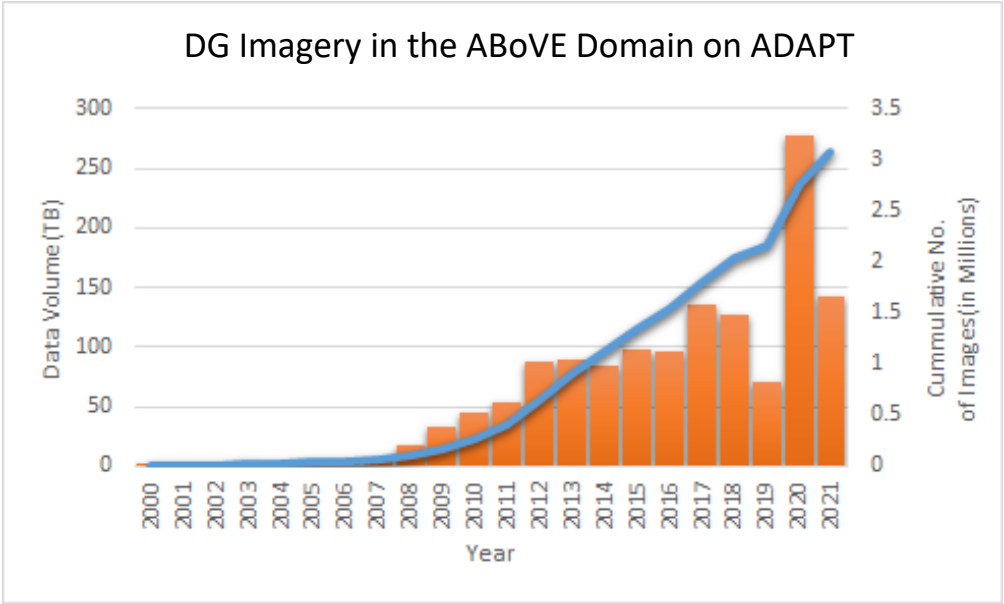

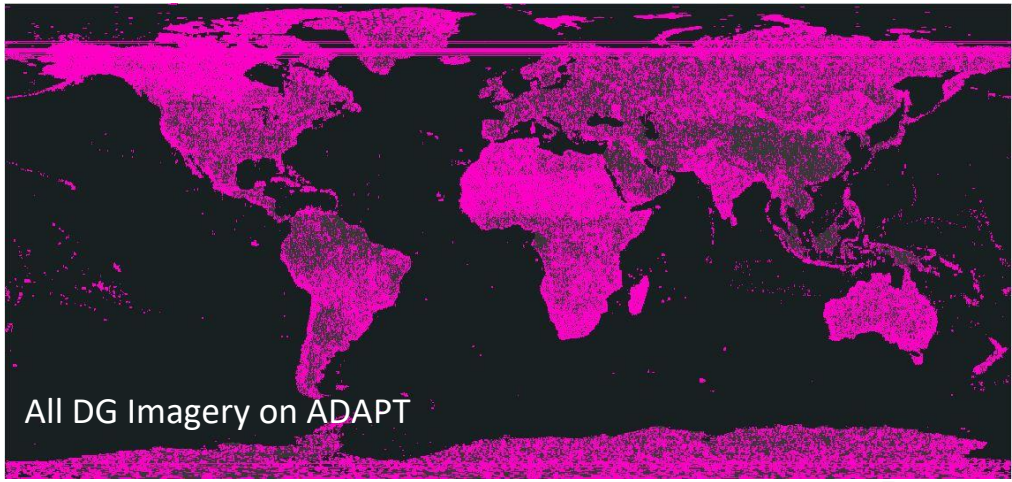

ABoVE-Archived Data from ORNL DAAC, NSIDC, and ASF are on ADAPT.

Location on ADAPT:

Linux:

```
-sh-4.2$ ls /css/above
AVIRIS_Analysis_Ready  datapool.asf.alaska.edu
daac.ornl.gov          n5eil01u.ecs.nsidc.org
```

Windows:

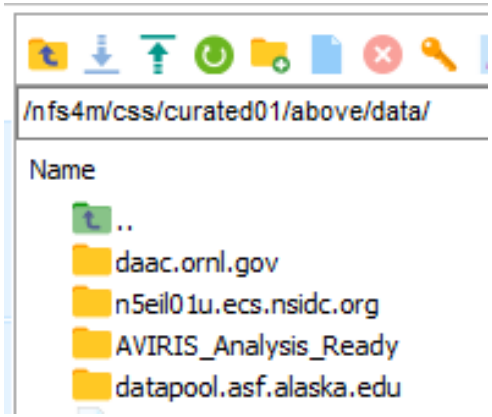

ESRI Capabilities on ADAPT

- Esri GIS Portal
  - https://maps.nccs.nasa.gov
  - 6 ArcGIS servers
    - Managed via the NCCS Windows environment
    - Direct to NCCS data (CSS, etc.)

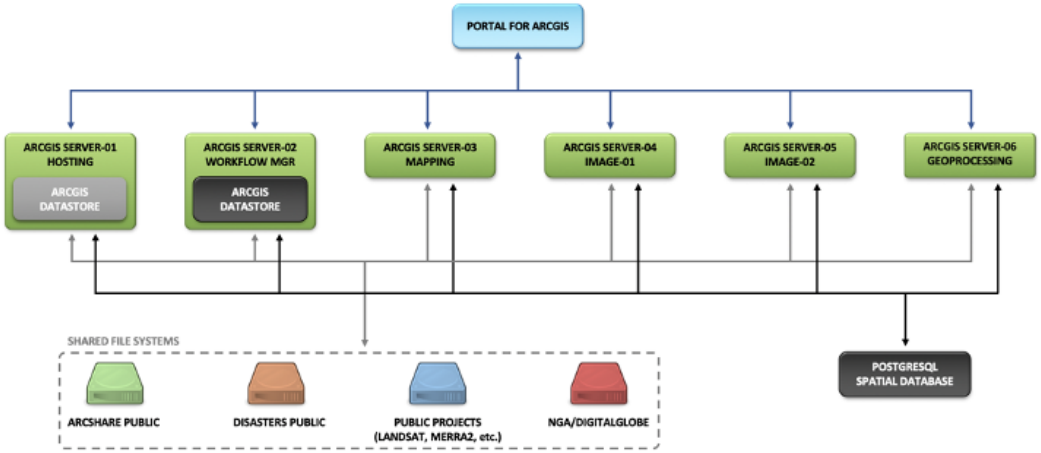

Discover

- ABoVE has an allocation on Discover
- Contact Liz if interested in using it

PRISM - GPU Cluster

- 22 nodes each with 4x NVIDIA V100 GPUs with 32 GB of VRAM
- 20 Intel CPU cores and 768 GB RAM per node
- Access requires an NCCS account

Jupyter Hub

- ABoVE has a Jupyter Hub instance
  - Currently only available for researchers on NASA VPN
  - https://www-proxy-dev.nccs.nasa.gov/jupyterhub-above/

- A second Jupyter Hub is available on PRISM
  - primary use is for researchers using GPUs, however also accessible on as-available basis
  - Does not require VPN.

ASC Updates

- ADAPT downtimes – Monthly, the Wednesday after the second Tuesday of the month.
  - October 13th
  - November 10th
- Moving to the Explore system - new hardware is coming in the next few months

ADAPT Issues? Contact support@nccs.nasa.gov

Announcements

- **Success stories** you would like to share? Email Liz Hoy elizabeth.hoy@nasa.gov
- How do I **cite the ASC in my publications**? Use language similar to:

> *"Resources supporting this work were provided by the NASA High-End Computing (HEC) Program through the NASA Center for Climate Simulation (NCCS) at Goddard Space Flight Center."*

Questions?

Tips and Tricks

Time for Matt to provide some tips and tricks