National Aeronautics and Space Administration

NASA Center for Climate Simulation (NCCS) User Forum

Laura Carriere, NCCS HPC Lead

9/14/23

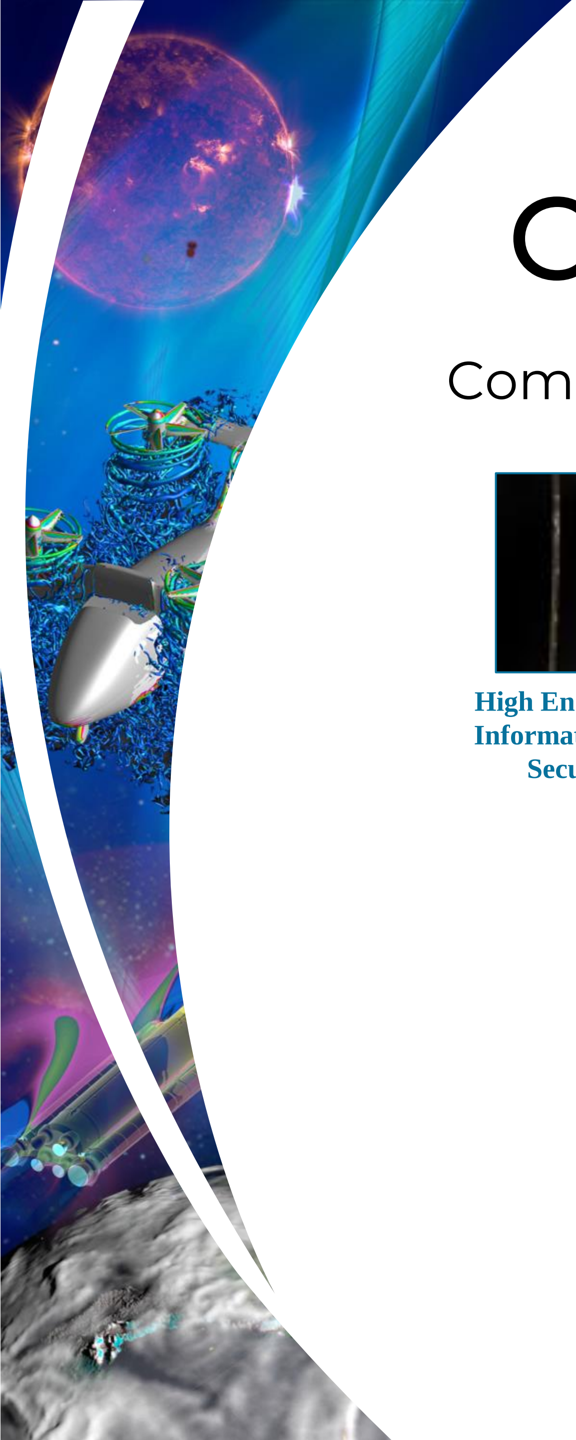

CISTO Organization

Computational and Information Sciences and Technology Office (CISTO)

High End Networks and Information Technology Security Group

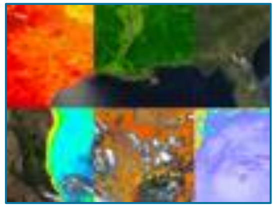

Data Science Group, Large Scale Analytics, Artificial Intelligence, and Machine Learning

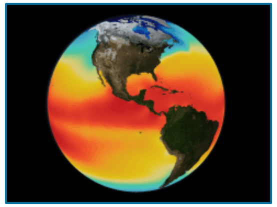

Scientific Visualization Studio (SVS) – *funded outside Scientific Computing*

NCCS – High Performance Computing (HPC)

Advanced Software Technology Group

Science Managed Cloud Environment (SMCE)

Discover Compute Resources

SCU14

Discover CPUs

- Traditional HPC
- 1,836 nodes, 80,000 Intel cores (today)
- 90,000 AMD Milan cores coming online this month, 90,000 more AMD Milan next summer
- IB and OPA Fabrics
- Parallel filesystem (Spectrum Scale)
  - 68 PB of storage, 110 PB on tape
- Slurm Scheduler
- Expansion occurs in scalable units

Discover GPUs

- CUDA and AI/ML
- 12 nodes
- AMD Rome CPU cores
- NVIDIA 4 x A100 GPU cores

Discover GPUs

SCU16 Compute Resources

SCU16 CPU

- Traditional HPC
- 676 nodes, 48 cores/node, 2.49 TF, 192 GB RAM/node, 4 GB RAM/core
- Intel Cascade Lake CPU cores
- IB Fabric
- SLES 12 SP5
- Replaced SCU11/12 in late 2021

SCU16 GPU

- CUDA and AI/ML
- 12 nodes
- AMD Rome CPU cores
- NVIDIA 4 x A100 GPU cores
- SLES 12 SP5

SCU17/18 Compute Resources

SCU17

SCU17

- Traditional HPC
- AMD cores – 90,112 Milan cores into operations later this month
- 704 nodes, 128 cores/node, 2.88 PF
- 512 GB RAM/node, 4 GB RAM/core
- IB Fabric
- SLES 15
- Benchmarks and early pilot users running now
- Operations later this month
- Replaced SCU10/13 (Haswell nodes)

SCU18 – Summer 2024

- Traditional HPC
- AMD cores – 90,112 Milan cores
- 704 nodes, 128 cores/node, 2.88 PF
- 512 GB RAM/node, 4 GB RAM/core
- IB Fabric
- SLES15
- Ordered
- Delivery May 2024
- Will replace SCU14/15 (Skylake nodes) – **upgrade challenges**

Discover Capacity Projection

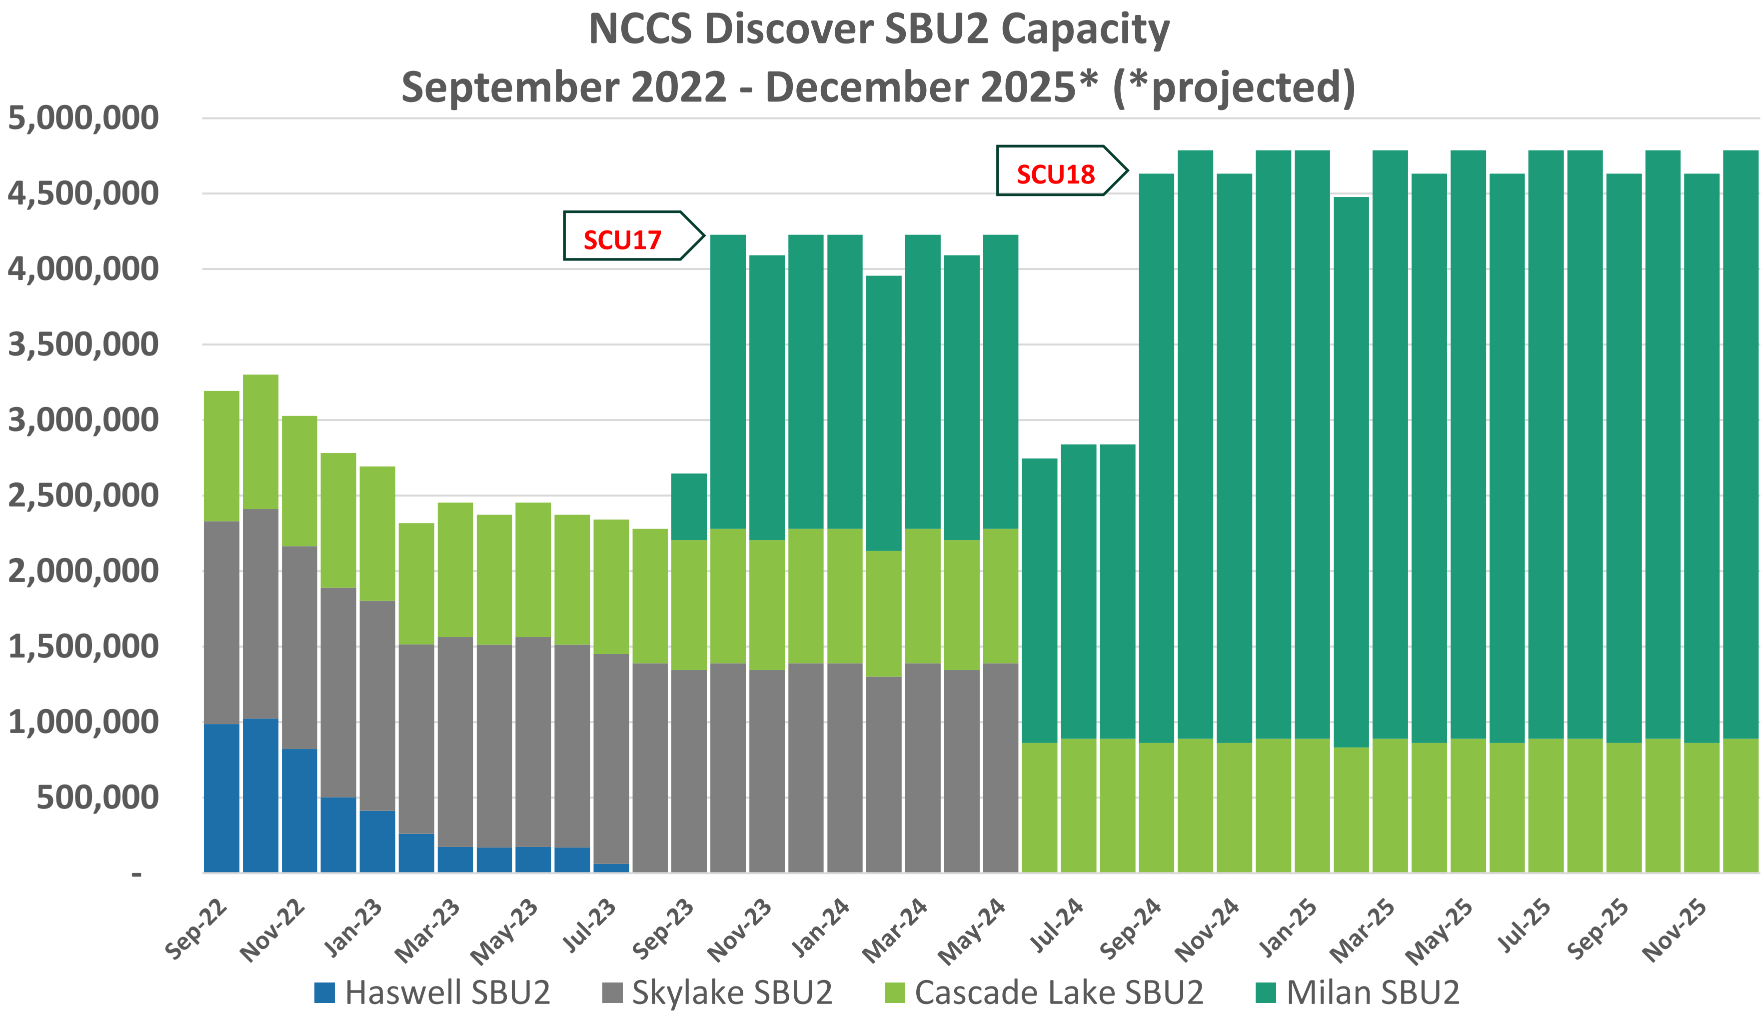

NASA High-End Computing

Historical Growth

Discover Base Unit (2006)

- 130 nodes, 4 cores/node
- 520 total cores
- 3.3 TF
- $515K/TF

Discover SCU17 (2023) – 1 chassis

- 4 nodes, 128 cores/node
- 512 total cores
- Chassis – 16.4 TF
- $4K/TF

Explore/ADAPT/Prism Resources

Explore (On-Prem Cloud)

- Appropriate for customized VMs, IaaS and PaaS
- 256 systems, 8,100 cores
  - Intel Broadwell, Haswell
  - AMD Naples, Rome
- Panasas Parallel filesystem, 7 PB of storage

Explore Container

Prism GPUs

- 22 systems with 4x NVIDIA V100
- 1 DGX system with 8x NVIDIA A100
- AI/ML software suites available

Prism GPUs

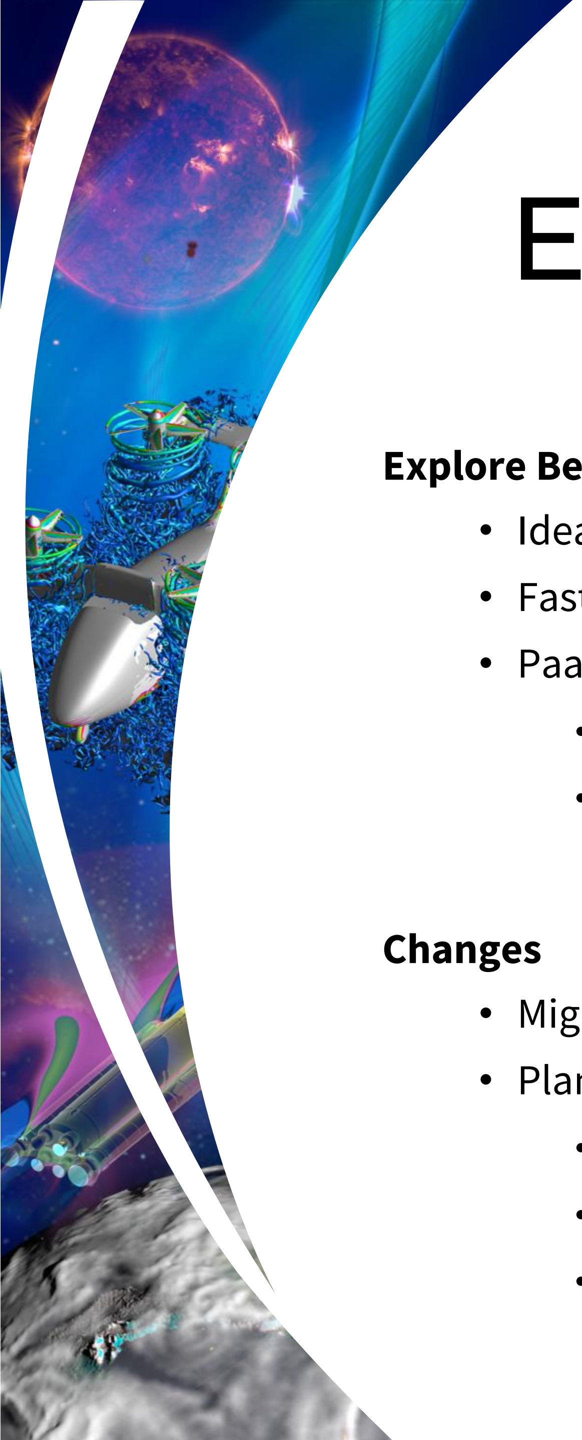

Explore Tenants and Changes

Explore Benefits

- Ideal for collaboration across a NASA team
- Fast access to hardware
- PaaS:
  - NCCS maintained OS
  - Access to NASA curated data products

Example Tenants

- ABoVE
- ICESat-2
- Roman ST
- Planetary Geodesy
- Astrophysics Anomaly Detection
- CALET (Calorimetric Electron Telescope)

Changes

- Migration to Panasas storage hardware for improved stability and performance
- Planned:
  - Improved patch installation process to reduce downtime during maintenance windows
  - 60-day limit on idle Prism /lscratch data
  - Move tenant VMs to Rocky 8

Future Plans

NIST SP 800-223 HPC System Reference Model

- Representative of Discover

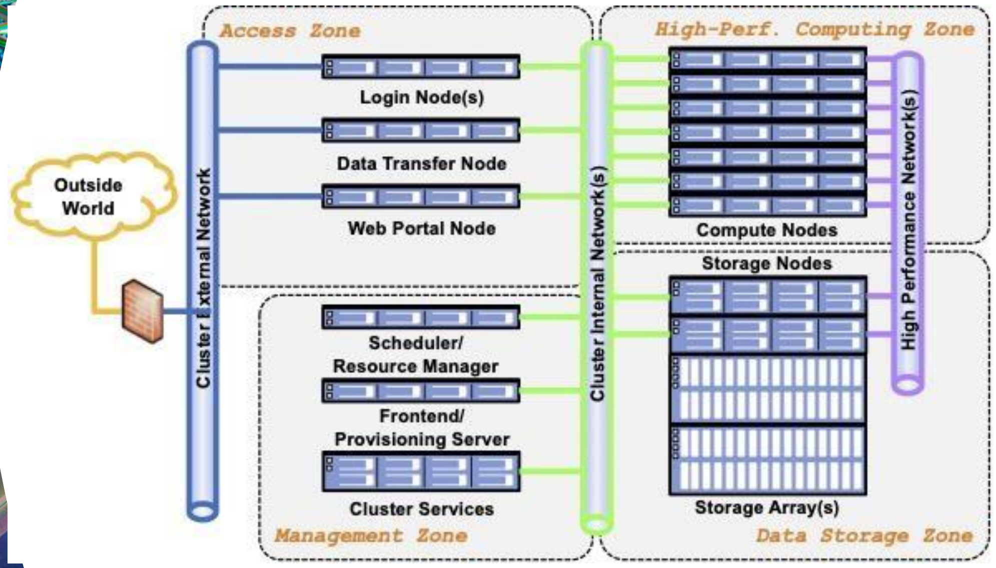

Credit: NIST SP 800-223

Next Generation

- Mix of HPC, AI/ML, customizable VMs, containers, Big Data, commercial cloud
- Challenges:
  - Ratio of CPU to GPU
  - Upgrades – power/cooling constraints
  - IB vs high speed ethernet
  - Tiered storage, no tape
  - Remove barriers to entry
  - Software installs – standardized for easy use, customizable for advanced users
  - Access to outside world for CI/CD, JupyterHub
  - Leverage Agency LDAP
  - Utilize CI/CD for provisioning, maintenance
  - Prepare for Zero Trust Architecture

Next Generation

- **SCU18 will be the last Scalable Compute Unit of Discover**
- Discover will continue to exist for 4-5 more years, with the SCU16/17/18 hardware
- FY24 funds will be used to:
  - Procure additional GPUs
  - Procure prototype hardware for the next generation system
  - Support more HPC in the Cloud

Strategic Compute Plans

- **“When everyone is super, no one will be”, The Incredibles**
- Discover: If every job has the highest priority, no job has the highest priority

- Each organization bears the responsibility for identifying the highest priority tasks within their organization
- Strategic plans would help identify that work and help us to help you get that work done, whether on Discover, at Ames, or in the Cloud
- We can help get your highest priority work through

Discover Queue Wait Times

The best way to help yourself, now and in the future is to make sure your requested wallclock time is as close as possible to your actual wallclock time. To do this:

1. Check your job summary report (slurm-<jobid>.out file, sent to your jobs standard output/error file
2. Compare actual wallclock time to the requested wallclock time
3. If actual is much lower than requested, revise your requested wallclock to be closer to the actual

Slurm will then be able to fit your jobs onto available nodes sooner. This will reduce your wait times for jobs to start.

Also – check your --constraint value to ensure you’re not excluding entire core families

Packable Concept

Example 10 core job:

Node	# Cores	Cores/Job	Idle Cores
Haswell	28	10	18
Skylake	40	10	30
Cascade Lake	46	10	36
Milan	128	10	118

- This isn’t going to work.
- The “Packable” QoS allows users to put multiple jobs on a single node to use the maximum cores
  - Currently this is constrained to the same user being able to run multiple jobs.
  - We need to move to a model where multiple users can “pack” multiple jobs on the same node.
  - To do this successfully, we need to ensure that each job cleans up after itself without affecting other users’ jobs

Packable Criteria and Plans

- Currently:
  - Jobs that use fewer half the minimum number of cores could be considered
    - < 20 for Skylake
    - < 23 for Cascade Lake
    - <64 for Milan nodes
  - Jobs must also use less than half the RAM on a given node
  - Contact Users Support at support@nccs.nasa.gov if you're are interested in trying this now
- In the near future (Fall 2023):
  - SCU17 (128 core Milan nodes) will be configured to allow packable by multiple users
- In the future (Winter 2024):
  - SCU14/15 and 16 will be configured to allow packable by multiple users

HPC in AWS

Discipline	Project	Job Count	Est. SBUs	Est. Cost
Earth Science	GMAO – Integrated Earth System Analysis Data Assimilation System	46	902	$1,700
Earth Science	GISS Global Gridded Crop Model Intercomparison	291	1,187	$1,000
Astrophysics	NICER	154	2,688	$3,200
Totals		**491**	**4,777**	**$5,900**

- Identified six projects
- Three were very successful, others had non-AWS/HPC challenges
- Pros: Short wait times to access fast cores, great support
- Cons: Support doesn't scale well yet, working on it
- Working on new partnerships with AWS and Azure

Criteria for HPC in AWS

- Easy to do:
  - Code development (small runs)
  - Well understood science (easy to estimate costs)
  - Reasonable data requirements
    - Ideal if data is already in AWS
    - Not more data than you want to copy in or out
  - Limited parallel I/O requirements due to cost of Lustre filesystem
    - Most code fits this description, even if it uses parallel I/O on Discover
- Also interested in trying jobs that don’t fit the above criteria
- Willingness to spend some time working with the SMCE HPC team to get set up (data, compilers, Slurm)
- If interested, send email

SCU17 Software Changes

- SCU17 will run SLES15
- As with previous OS upgrades, we will be limiting the number of software modules that we will be carrying forward, unless there's a very specific requirement
- SLES15 does not have **Python 2, only Python 3** and we will not be installing it. The Agency considers it to be insecure.
- You must migrate to Python 3 to run on SCU17
- Please watch for:
  - Tech Talks, e.g., Best practices for effective use of Discover, resilient scripts, ability to restart code
  - ASTG Python Courses (and Fortran)
  - NCCS Nuggets, e.g., Requested wall clock times

GPU Allocations

- Planning to purchase additional GPU resources in FY24
  - Will likely procure 1 rack of NVIDIA 4 x A100s
  - Still determining placement (Prism, Discover, split between the two)
- Starting this year, users can make specific GPU allocation requests through RMS
  - Not a requirement this year
  - Will become one for future systems
  - Will still have discretionary project ids for new users
- Purpose is to help us determine future requirements
- Poll

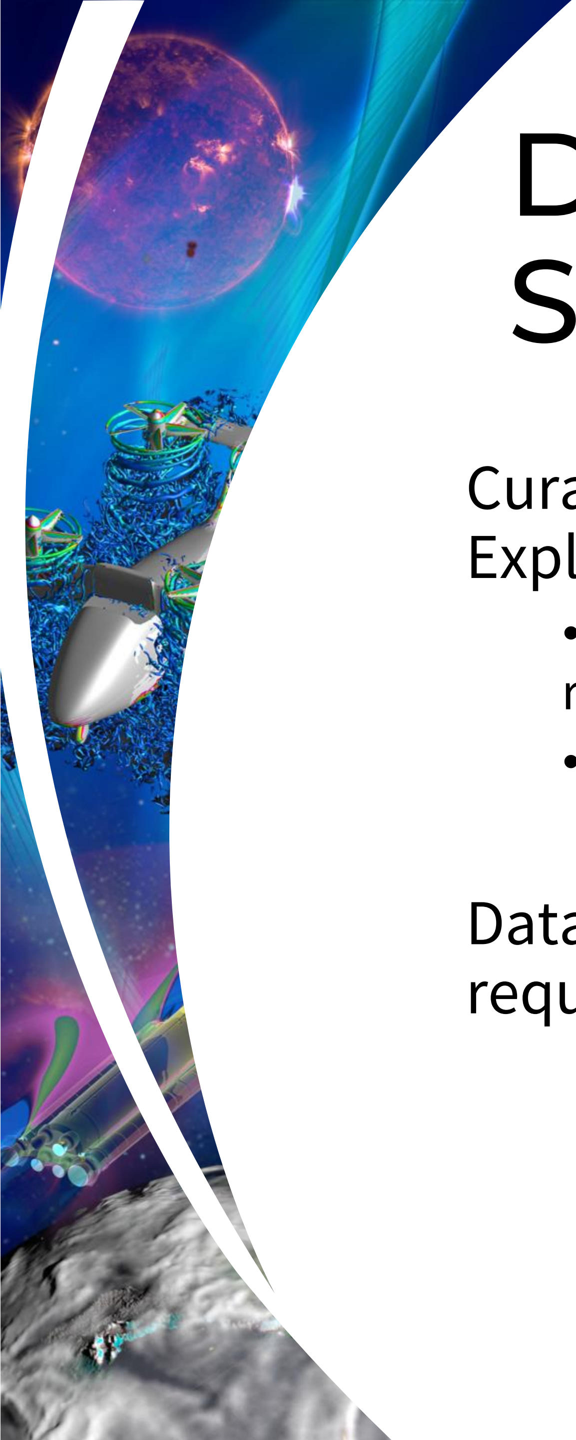

Data – Centralized Storage System (CSS)

Curated NASA data products examples available from Discover, Explore, and Prism (45 PB of data):

- CMIP5/6, Landsat, MODIS, MERRA2, NEX GDDP, ICESat-2, NGA (very high-resolution imagery)
- Search on Data Collections on the NCCS website

Data Management Plans are used to manage large (100TB+) storage requirements

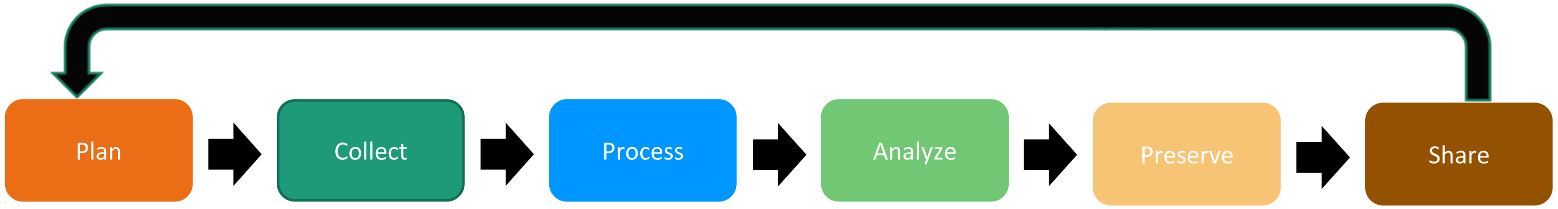

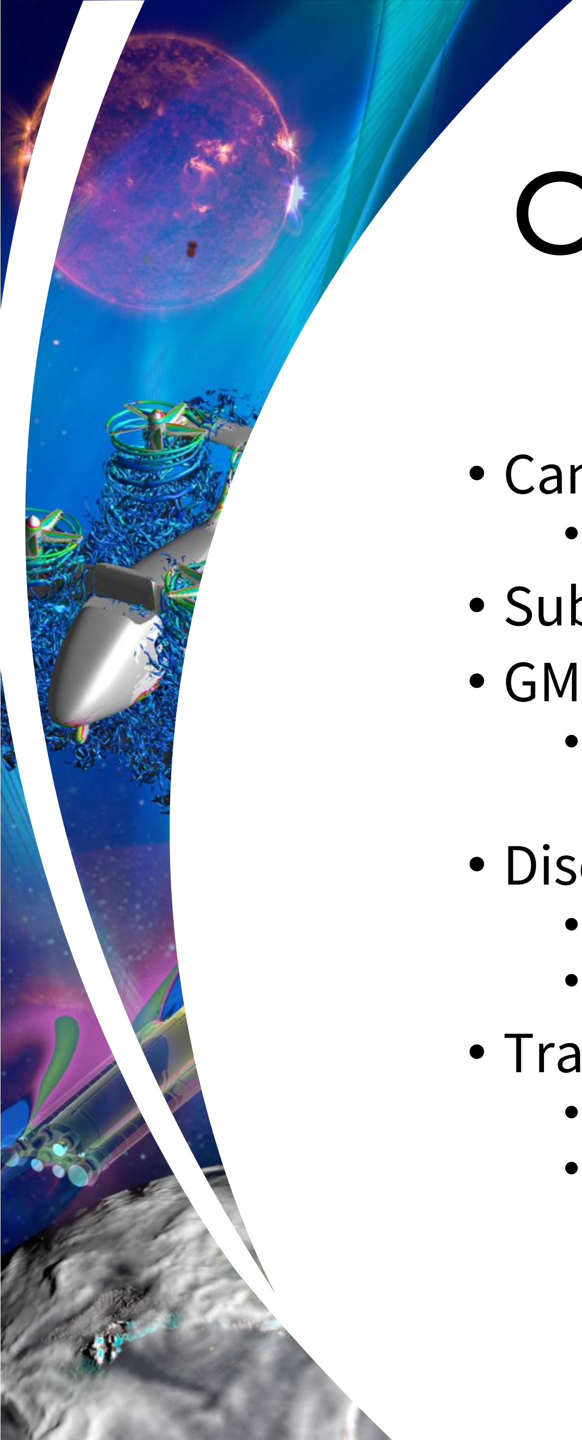

Changes to Storage

- Campaign storage (CSS) available to SCU16 and SCU17
  - Read/write from login nodes
- Submit jobs with constraint=cssro for access from compute nodes
- GMAO input data is now located on disk on Discover:
  - /discover/nobackup/projects/gmao/input/gmao_ops/ops
    - Either flk or reanalysis, consult the index file obsys.rc
- Discover data migration
  - GPFS version change from v4 to v5 requires data copies
  - More storage (18PB) is coming online later this fall
- Transactional storage
  - 1 PB of NVME, 4x faster than GPFS
  - Access via $TSE_TMPDIR, deleted at end of job

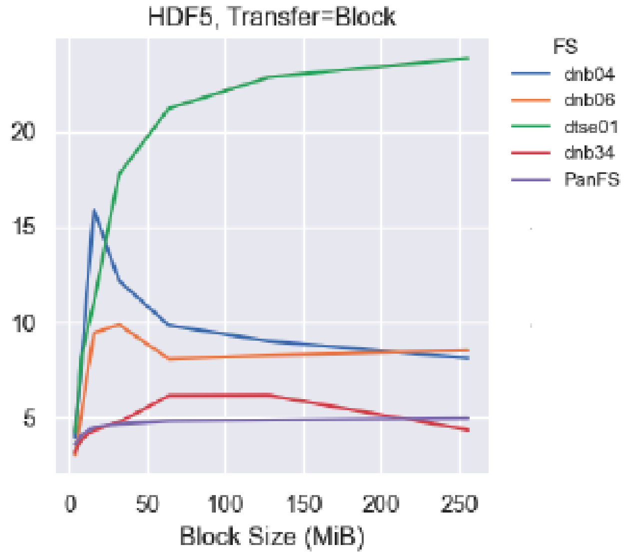

NCCS Storage Policy

- Will be posting a summary of the policy, including the default storage allocations across all NCCS environments
- NCCS provides, maintains, and monitors storage
- NCCS Users:
  - Provide DMPs where required and apply data management life cycle best practices
  - Inform their PIs and the NCCS when they leave. PIs/Sponsors must ensure that ownership is transferred to another person within the project or to another project or indicate that it may be removed
  - **Orphaned data**: If data remains unclaimed and/or the NCCS receives no direction from the PI, it will be made inaccessible and will not be retained for longer than 2 years. If storage resources become constrained, data may be deleted sooner

Changes to Mass Storage

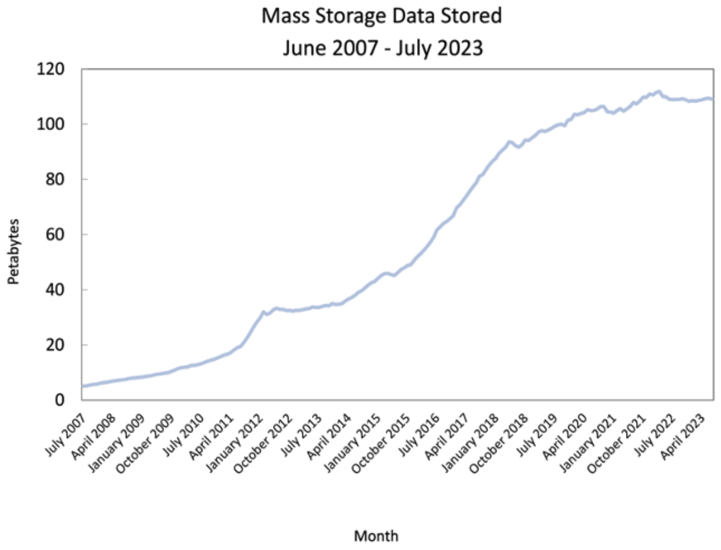

- Moving to read-only
- Only 1 major project still writing to MSS
- Other workflows moved to disk
  - Discover for input and intermediate
  - CSS for final
- Developing processes and costs for storing data at NAS that isn't input, intermediate, or final
  - Data must have a DMP to be considered for NAS storage
- Working with NASA on a Climate Model DAAC concept
- Starting to work on extracting data so that it can be shut down

JupyterHub

- Discover (CPU and GPU), Prism, Explore by tenant
- AI/ML and CUDA
- Standardized for quick start up
- Customizable for advanced use
- Example Notebooks are available
- All Discover users have access to Discover JupyterHub
- All Explore users have access to Prism JupyterHub

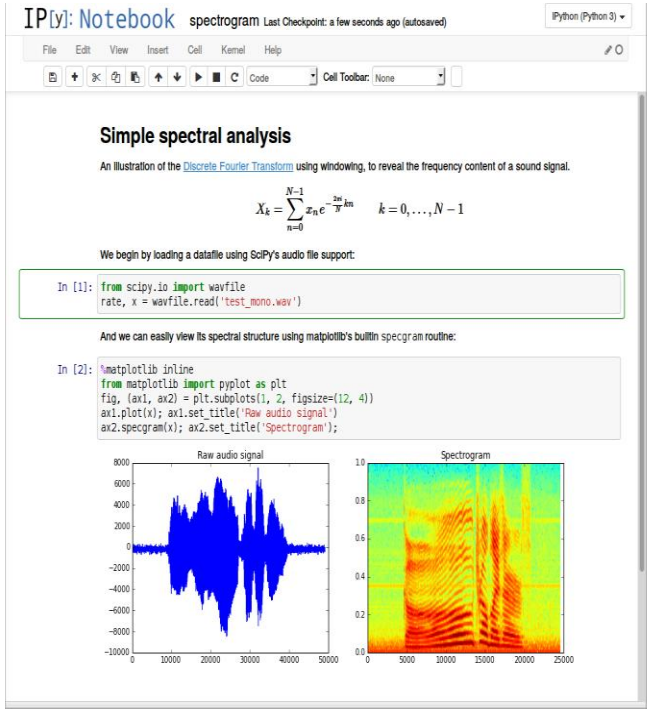

Gitlab

- Available to NCCS projects
- To request access:
  - NAMS request NCCS Gitlab Service # 257778
- Code in Gitlab is backed up

- More information:
  - https://www.nccs.nasa.gov/nccs-users/instructional/gitlab
  - Or go to nccs.nasa.gov and search on Gitlab

Support Notes

- New users:
  - Email support@nccs.nasa.gov for discretionary
  - RMS for allocations – request.hec.nasa.gov
- NAMS Account Request Process for NCCS Accounts
  - GSFC NASA Center for Climate Simulation #257778
- Ticket Guidelines
  - NCCS Website - Contact-us/ticket-guide
- Intel/NVIDIA Support
  - Can organize consultations with both (check NVIDIA)

Questions?

Software - Explore

- Standardized:
  - ML and MLGPU Kernels
    - Traditional AI/ML tools: Tensorflow, Keras, Pytorch, Scikit Learn
  - ILab Kernel
    - Intended for AI/ML workflows specific to earth and environmental science
  - Sci2 Kernel
    - Large collection of scientific applications, including support for climate modeling
  - EViz consists of two Python-based tools, for visualizing a wide variety of data sources.
    - eViz is a CLI-driven static plotting tool configurable with YAML files
    - iViz is a Jupyter Notebook-based tool with a Bokeh backend for providing an interactive exploration of data sources.
  - Octave, R
- Customizable by tenant VMs

Software - Discover

- CI/CD Pipelines
  - Gitlab, Singularity, CharlieCloud
- NCCS Managed
  - Modules
  - COTS: Matlab, IDL, ARM Forge, Totalview Compilers (Intel, PGI, NAG Fortran)
  - Open Source: Gcc, Python, HDF, NetCDF, Octave, Julia, Tau, CUDA
  - Machine Learning: TensorFlow, PyTorch, others
- Customer : GEOS, ModelE, LDAS, NU-WRF
- User Managed
  - Conda, Spack
- Visualization
  - JupyterHub, Visit, Panoply, ImageMagick