NASA SMCE & Microsoft Webinar

Daniel Duffy
CISTO Chief/SMCE Manager
Code 606, GSFC
daniel.q.duffy@nasa.gov

February 5th, 2025

Objectives

1. Provide guidance on accessing and utilizing Microsoft Azure services through NASA's Science Managed Cloud Environment (SMCE)
2. Introduce Microsoft Azure's technologies, resources, and training materials to get started.
3. Demonstrate use cases of Azure's cloud and AI capabilities applied to NASA science.

Agenda

Time	Session	Speaker
1:00PM	Welcome: Agenda and objectives	
10 mins.	**NASA's SMCE Overview** Introduction to SMCE capabilities Onboarding process for Microsoft Azure	**Dan Duffy** NASA
10 mins.	**Introduction to Microsoft Azure** Overview of Azure technologies for science including Planetary Computer, Azure OpenAI, compute, and others Available resources and trainings	**Juan Carlos López** Microsoft
15 mins.	**NASA Use Case: HPC** Leveraging Microsoft Azure High Performance Compute (HPC) to accelerate NASA science from months to weeks	**Dan'l Pierce** NASA SMCE
15 mins.	**NASA Use Case: AI & Data** NASA Earth Copilot: leveraging AI and Large Language Models to make NASA's Earth Science data and insights more accessible	**Muthukumaran Ramasubramanian** NASA IMPACT
5 mins.	Closing Recap, resources, and next steps	**Juan Carlos López** Microsoft

SCIENCE CLOUD
NASA

NASA

Accelerating Science Through NASA Cloud Computing

NASA SMCE/Microsoft Webinar
05 February 2025

Computational and Information Sciences and Technology Office (CISTO – Code 606)

Vision: Accelerating science through high-end information technology solutions; from bits on the wire through understanding.

High End Networks and Information Technology (HECN)

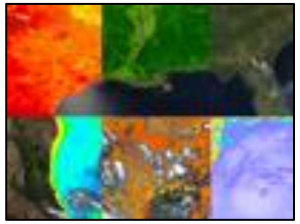

Data Science Group (DSG)
Artificial Intelligence, and Machine Learning

High Performance Computing (HPC) on-premises compute

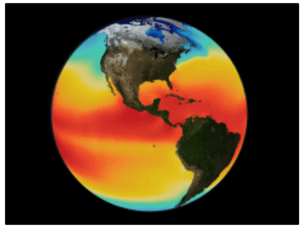

Scientific Visualization Studio (SVS)

Advanced Science Applications and Technology (ASTG)

SMD Cloud Infrastructure Project (SCIP) and the Science Managed Cloud Environment (SMCE) – AWS Commercial Cloud Resources

https://science.gsfc.nasa.gov/cisto/

The Evolution of Cloud Computing Services for Science

2015

The AIST Managed Cloud Environment (AMCE) Formed

- Incorporated NASA security, technical, and financial standards into cloud computing
- Fewer than 10 projects
- A single shared AWS Account

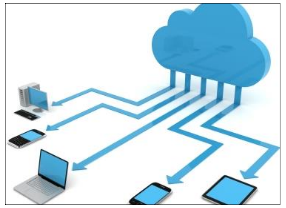

2018

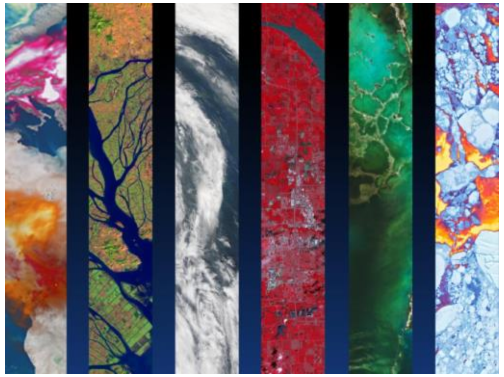

The Science Managed Cloud Environment (SMCE) Formed

- First major project: Earth Information System
- Migrated HPC workloads to the cloud
- Offers Azure and AWS capabilities
- Supports all Science Mission Directorates
- Now serves almost 100 projects

2024

The Science Cloud Infrastructure Project (SCIP) Began

- Aims to unify SMD cloud infrastructure into a single environment
- Will modernize and streamline NASA cloud computing efforts

The Science Managed Cloud Environment (SMCE)

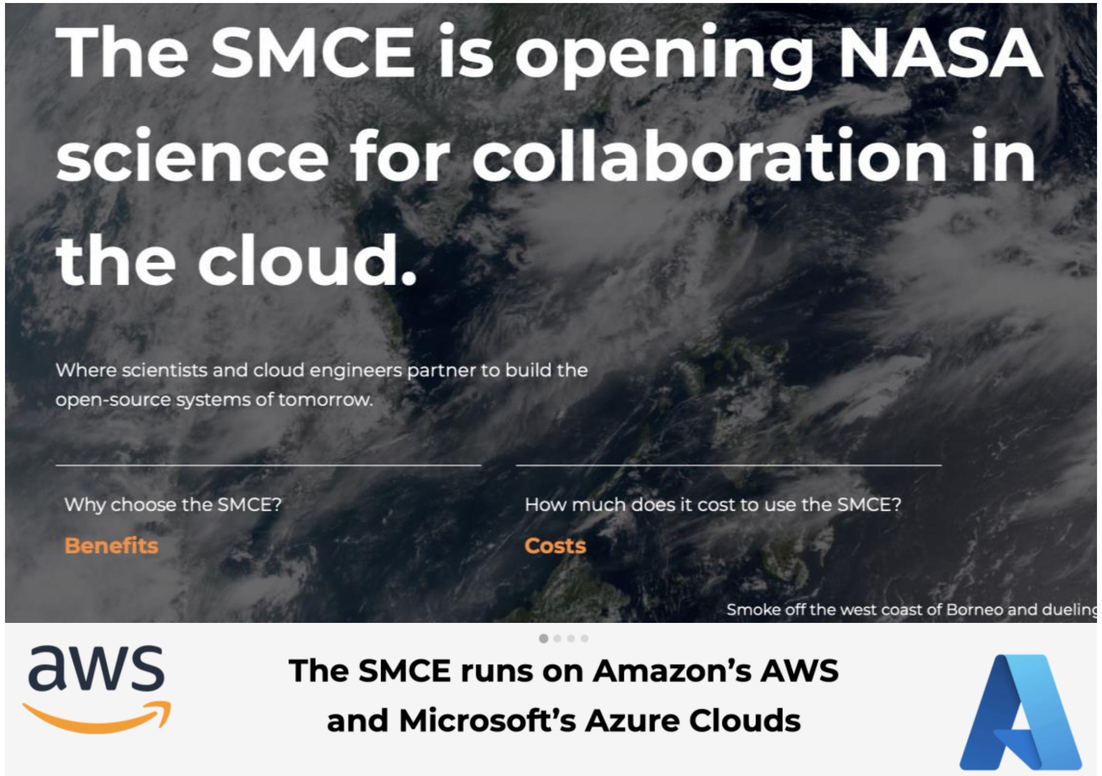

The SMCE offers:

- Direct access to Amazon's AWS and/or Microsoft's Azure
- Ready to go cloud-native Science Analytics Platform
- Open Science Studio (OSS), a scalable JupyterHub environment with customizations for science
- Rapid access to emerging cloud computing capabilities and services, such as Microsoft's OpenAI
- Secure collaborative environment for research, development, prototyping, open-science, and training
- Parallel clusters for HPC
- Gitlab repository and CI/CD capabilities
- FISMA-low environment
- Consulting, training, and support

Why HPC in the cloud?

Cloud-Based High-Performance Computing can eliminate constraints of the on-premises systems

- Significantly reducing wait time
  - Reduces limits on the number of jobs that can be run simultaneously
  - No shared queue
- Flexibility of environments and software stacks
- Access to the fastest processors available and to processors which match the application requirements

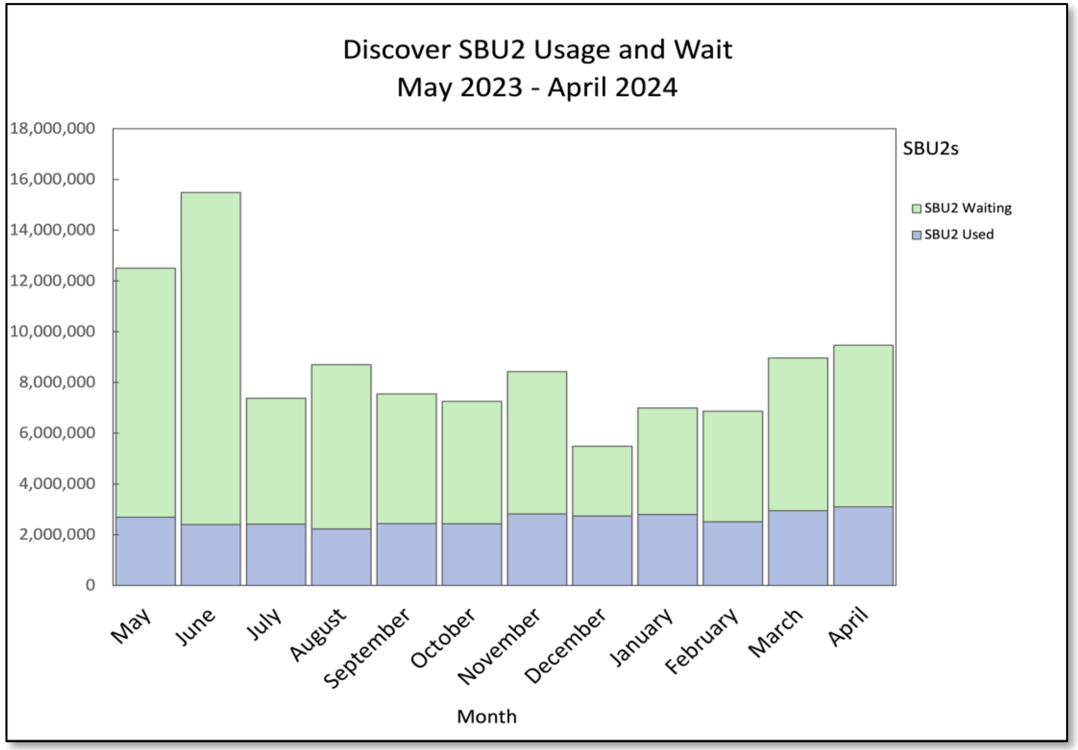

NCCS (on-premises) effective measure of the wait times for jobs throughout the month. It is not uncommon for the on-premises Discover HPC system to have 2x to 3x the requirements waiting versus executed over the course of the month.

Why Choose the SMCE or the Science Cloud?

"But the cloud is much more expensive than on-premises! Why would I ever use it?"

- Yes, it is. Costs have dropped significantly over the years, but it is still cheaper to do things on premises.

"You are not being very convincing!"

- If you can't get your work done effectively using on-premises systems, then cloud computing can be an alternative. This happens because of security requirements, the heavy demand, and even the policies that govern on-premises systems.
- Or what if you wanted access to some emerging technology such as the latest processors, AI/ML, or potentially even quantum computing? How long would you have to wait to get that on-premises? You will have the quickest access to these technologies by using cloud computing versus waiting for months or even longer than a year to get access on-premises.

"Okay, I can see some of that. But how could that help me?"

- What if we could significantly accelerate your research by rapidly creating a cloud environment specific to your requirements? We will show you an example of that.

"Fine, but who is going to pay for it?"

- The project pays for the resources used in the cloud; no overhead for the SMCE or the Science Cloud team.
- However, there are ways to get funding for cloud resources, including from High End Computing, funded research projects, and even some seed funding. Let's talk!

How To Access the SMCE

How to request access and start a project:

- NASA Principal Investigators (PIs) interested in establishing a project within the SMCE should contact the SMCE Team at

  smce-admin@lists.nasa.gov

  The SMCE Management and User Liaison will reach out to discuss requirements, solutions, cost estimates, and plans to move forward.

For more information…

- If you are interested in the SMCE or have further questions, please contact the SMCE team at

  smce-admin@lists.nasa.gov

Thank you to all the Team Members

SMCE Team	Science Cloud Formulation Team

SMCE Team

- Project Lead: Daniel Duffy
- Lead Architect: Ramon Ramirez-Linan
- Operations Lead: Aaron Skolnick
- User Liaison: Dan'l Pierce
- Dev/Ops: Jeremy Raupp
- HPC Dev/Ops: Dorian Crockrell
- HPC Special Projects: Hoot Thompson
- Support Engineer: Johana Chazaro Cortes
- Support Engineer: Sarah Au
- ISO: Bob Peirce
- Lead ISSO: Jay Ellis
- Dev/Sec/Opps: Andre Avelino Paniagua

→

Science Cloud Formulation Team

- Project Lead: Daniel Duffy
- Project Manager: Ruma Das
- Technical Lead: Garrison Vaughan
- Security SME: Jim Wiedman
- Comms Lead: Audrey McQuagge
- SMCE Product Owner: Johana Chazaro Cortes
- Data Divisions Product Owner: Shanna Sampson
- Cost and Req's Analysis: Aerospace
- Lead Architect: Kevin Mentzer
- Research and Innovation Lead: Ramon Ramirez-Linan
- Dev/Sec/Ops Engineer: Kat Morgan
- Dev/Sec/Ops Engineer: Kingdon Barrett
- Lead ISSO: Tara Leonard
- Security Engineer: Jasaun Neff

Note that the implementation/operations team for the Science Cloud is still being formed to include other members of the SMCE team and also the Earth Data Cloud infrastructure and security teams.

NASA Science Cloud

ACCELERATING SCIENCE BY REDUCING BARRIERS TO THE UTILIZATION OF COMMERCIAL CLOUD SERVICES

NASA

Microsoft Azure

Cloud powered innovation on and off the planet

Juan Carlos López
Former NASA Engineer
Azure Specialist, NASA
juanlopez@microsoft.com

Agenda:

Technology

People

Skilling

Microsoft Global Network

Hyperscale Cloud and Network

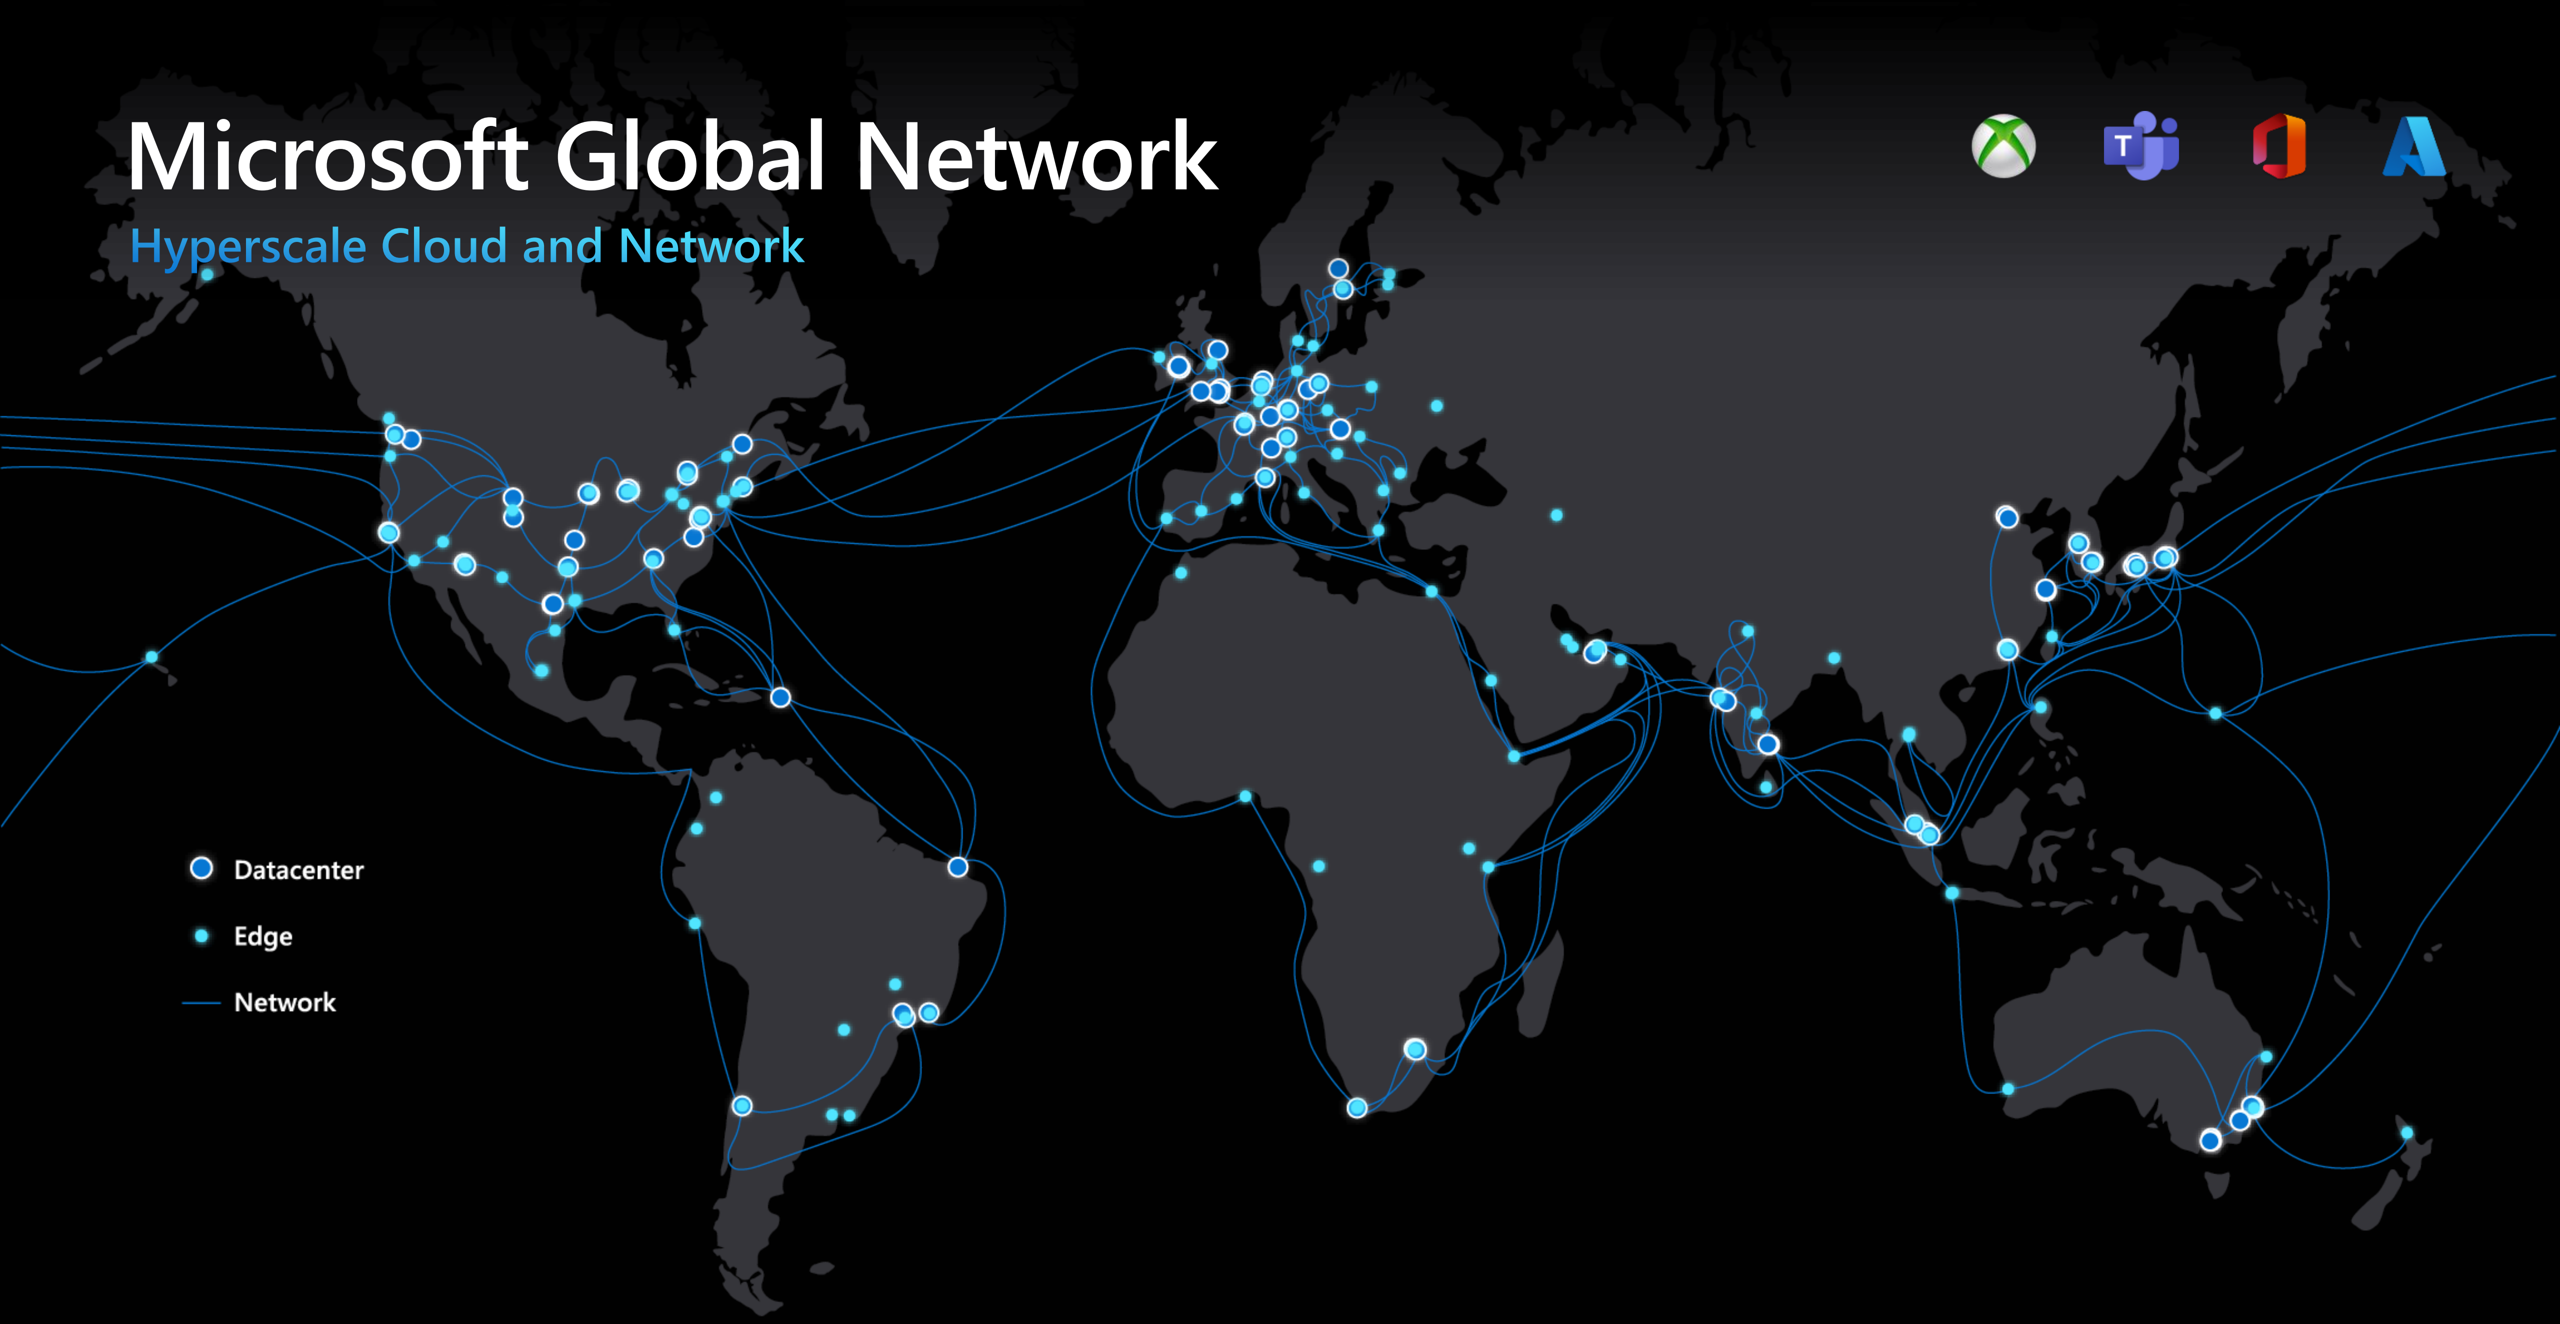

Microsoft

NDA Content

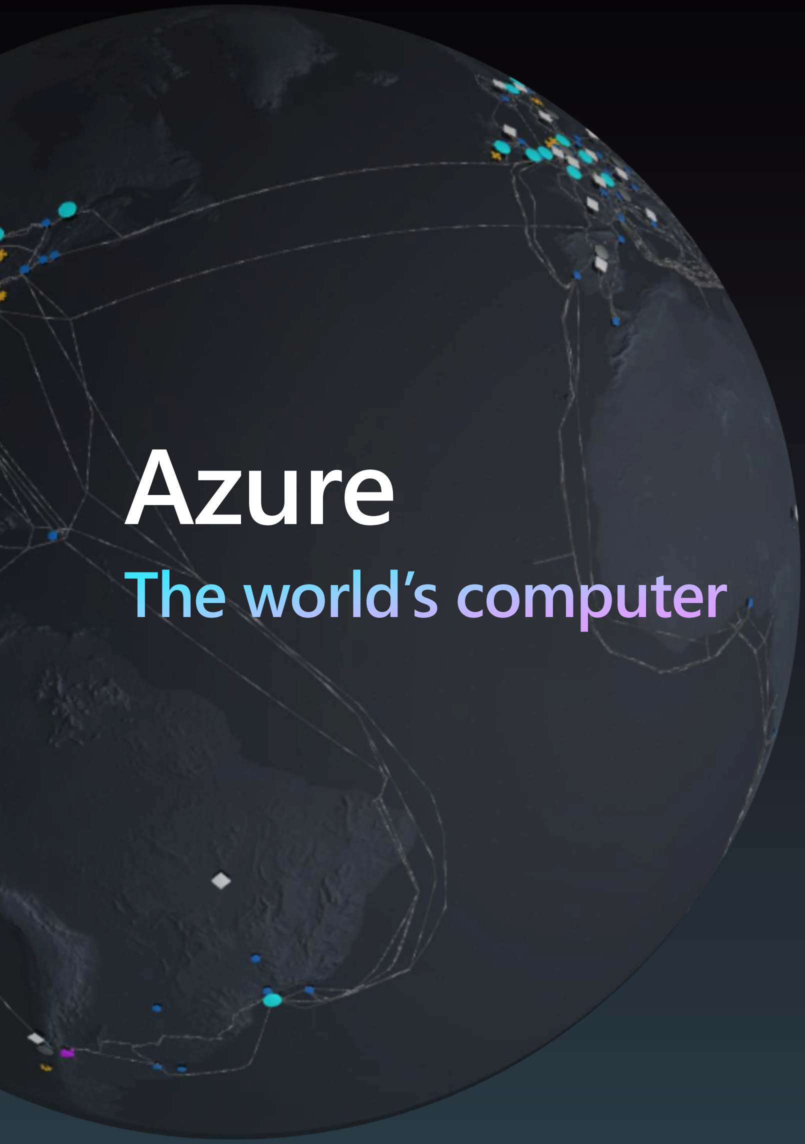

Azure

The world's computer

1 Intelligent data and AI

2 Services for all workloads

3 DevOps tools and experiences

4 Security, identity, and governance

Azure Arc

bridging services/apps to continuous computing fabric

5 Continuous cloud-to-edge computing fabric

Global infrastructure

Enterprise core

Operational edge

200 + Azure services available

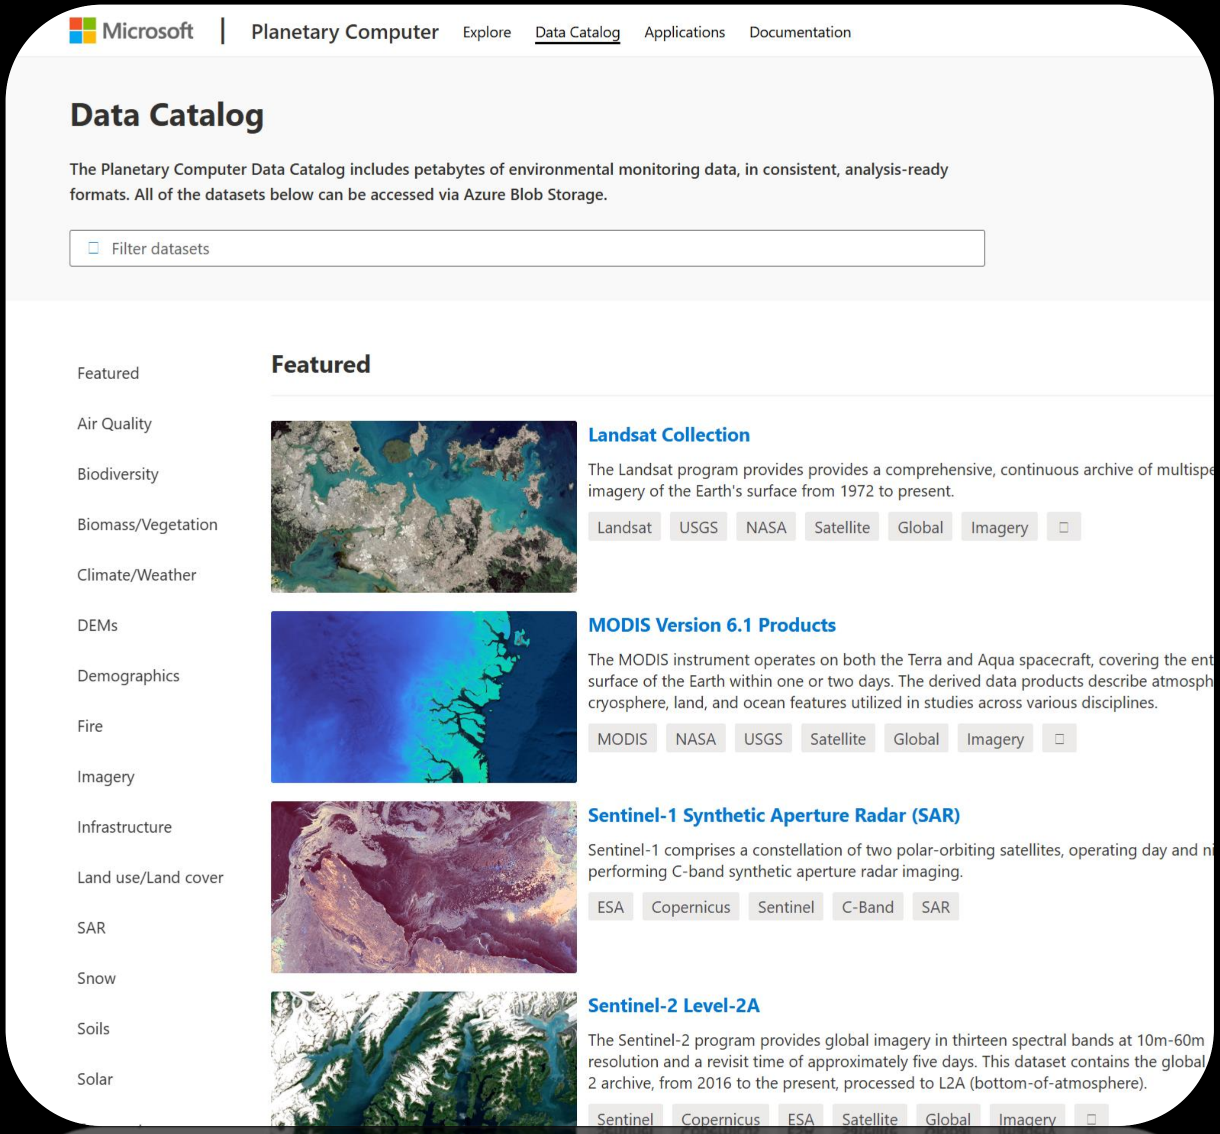

Microsoft

Planetary Computer

- Data Catalog, including:
  - Landsat
  - Sentinel-1
  - HLS *(coming soon)*
- API
- Applications

https://planetarycomputer.microsoft.com/

Azure AI Foundry

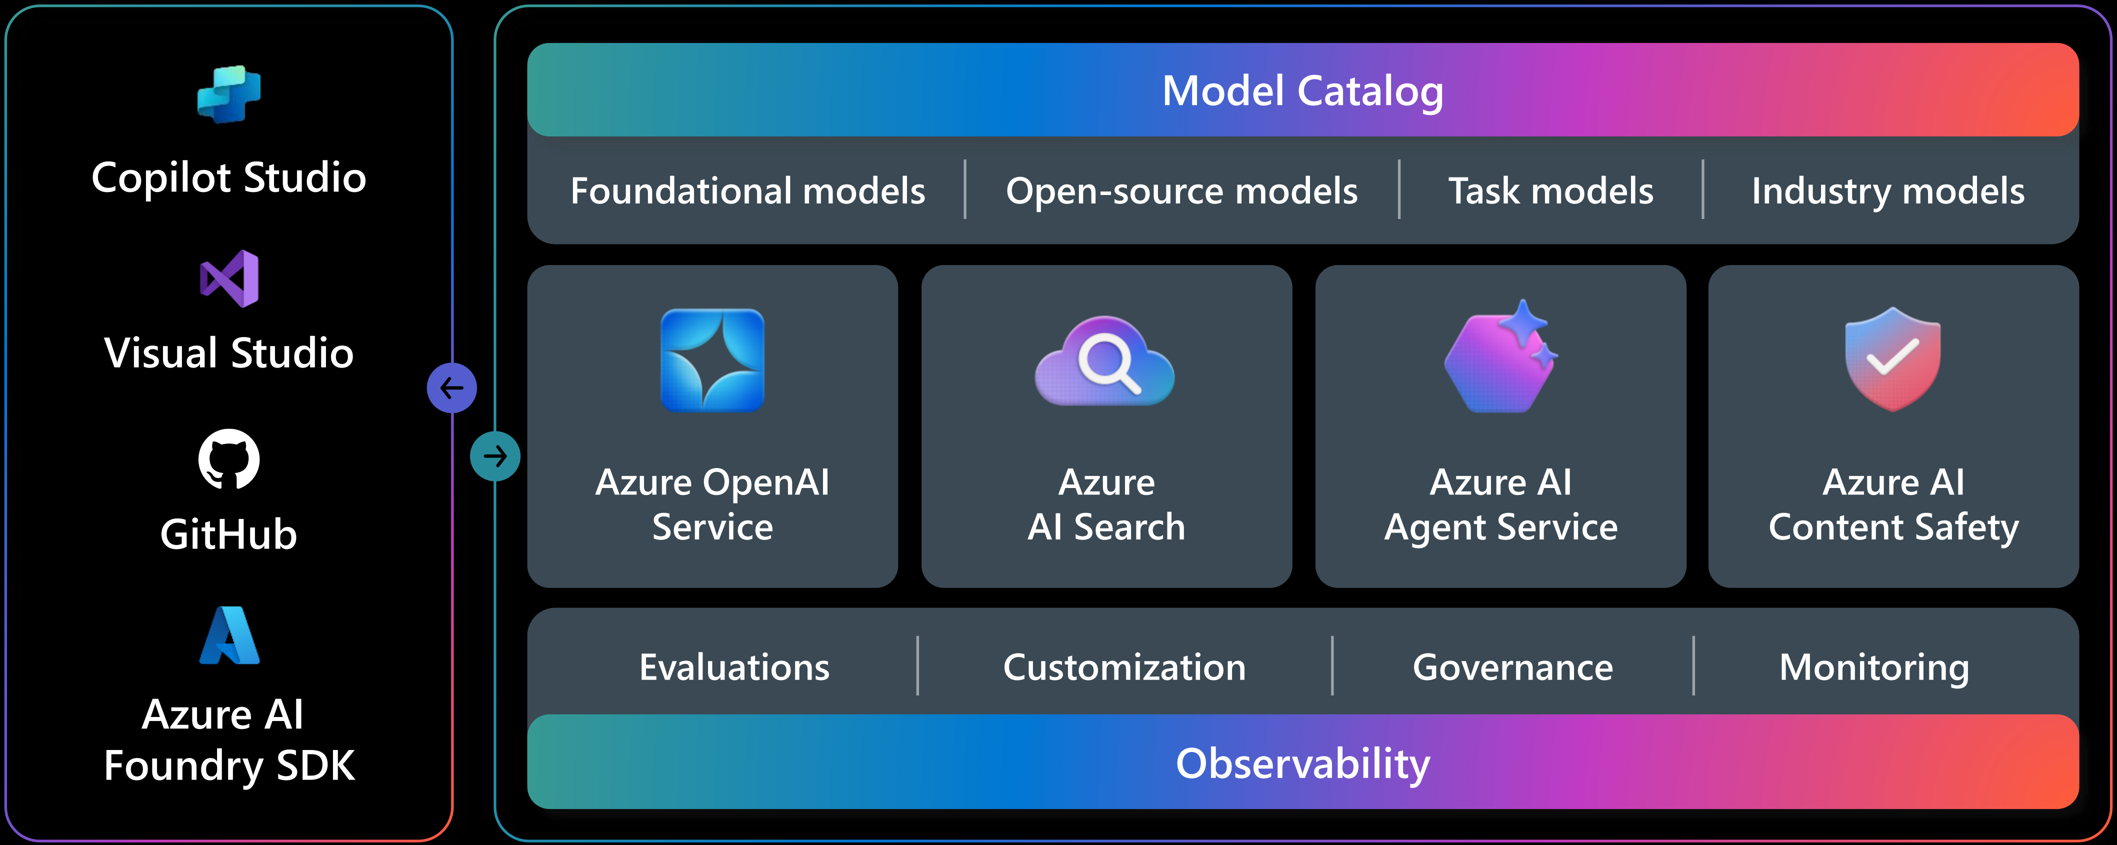

Azure High-Performance Compute

Accelerating Science

Purpose-built supercomputing

Supercomputing platform purpose-built for AI and HPC workloads.

Portfolio of virtual machines to meet any needs, with virtually unlimited scalability.

+

Intelligence services

Full range of AI and analytic services to fuel innovation.

Deliver AI, machine learning, deep learning and business insights securely, at scale.

+

Fast, secure networking

Fast InfiniBand inter-connects and as edge-to-cloud connectivity with 400 GB/s InfiniBand® and Azure ExpressRoute.

+

High-performing storage

High performing, secure and fully integrated storage options to support simple-to-complex needs.

+

Workload services

End-to-end workflow agility with integrated familiar tools and processes.

Support for industry-leading models, frameworks, containers…

NASA

Microsoft Azure

Applying for a NASA grant?
Starting a new science project?

We are here to help!

- Project Formulation
- Solution Architectures
- Cost estimates
- Prototyping

Juan Carlos López
Azure Cloud Specialist
Former NASA Engineer

Email: juanlopez@microsoft.com

Jack Wade
Principal Technical Specialist
Azure Infra

Paul Lizer
Principal Technical Specialist
Azure Data & AI

Jocelynn Hartwig
Senior Technical Specialist
Azure Data & AI

Adam Pavlik
Senior Technical Specialist
Azure Infra

Mahdi Khan
Technical Specialist
Azure Infra

Support Team

Already an Azure user via SMCE? **You have a group dedicated to your success!**

- Support tickets
- Expert Advice
- Dedicated Support Engineers
- Proactive engagements

Customer Success

Bringing right resources at the right time to develop and deploy solutions on Azure

Gwen Sedler

Samantha Falcucci

Tara Marchi

Services Executive

Incidents Manager

Reactive support case management

Saul Puente

Dan Biscup

AI Cloud Solution Architect

Dedicated Engineers

Your hands-on technical experts throughout your cloud journey, covering Azure infrastructure, identity, security, data, AI/ML, etc.

Angel Amaral

Platform Cloud Solution Architect

Seth Rimmer

Azure Infra Cloud Solution Architect

Michael Crane

Identity & Cybersecurity Cloud Solution Architect

Enterprise Skilling Initiative (ESI)

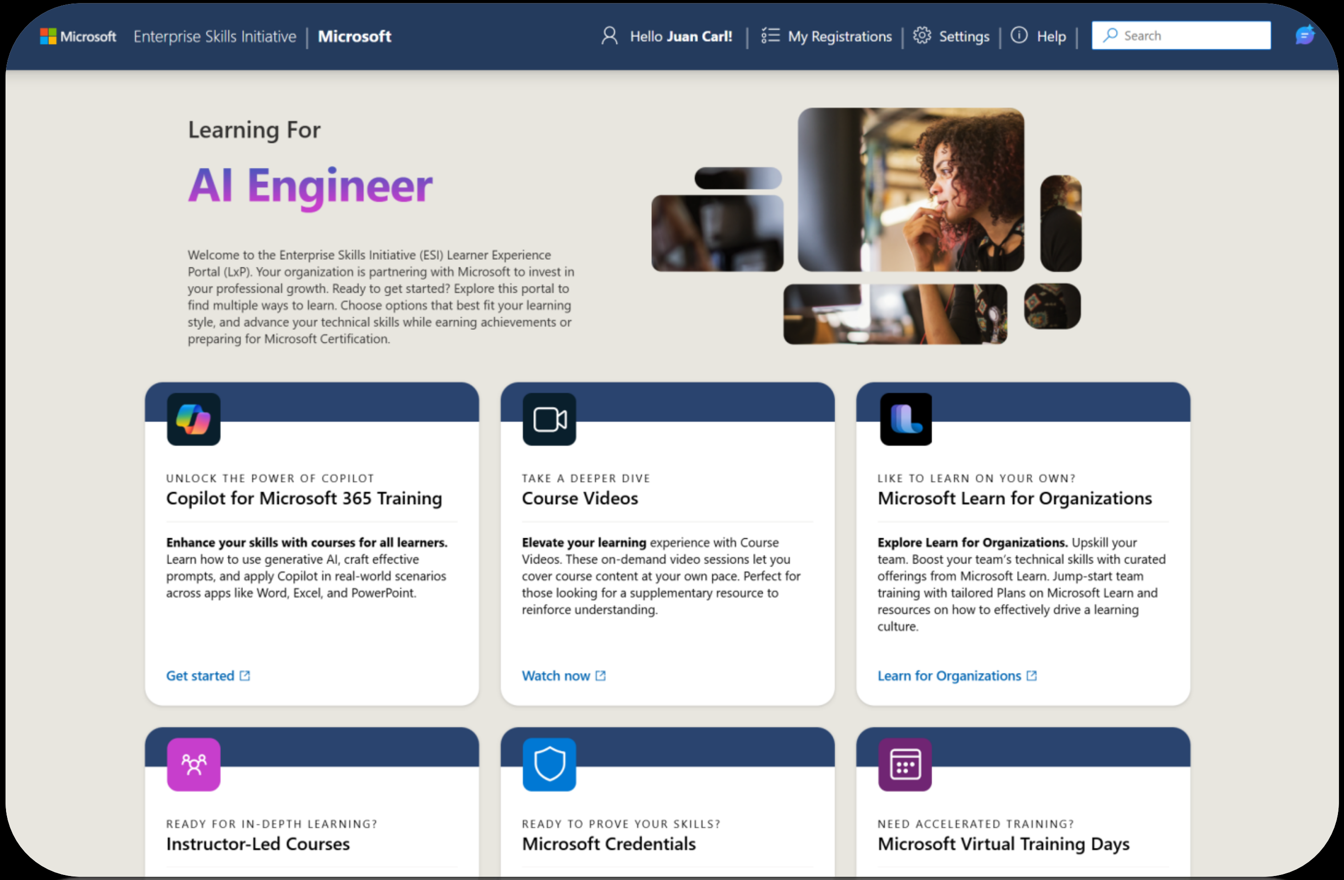

- Hundreds of FREE trainings available to NASA
- Create an account with your NASA.gov e-mail
- Discounts on certification exams!

https://esi.microsoft.com/

Questions?

Juan Carlos López
Azure Specialist
Former NASA Engineer
juanlopez@microsoft.com

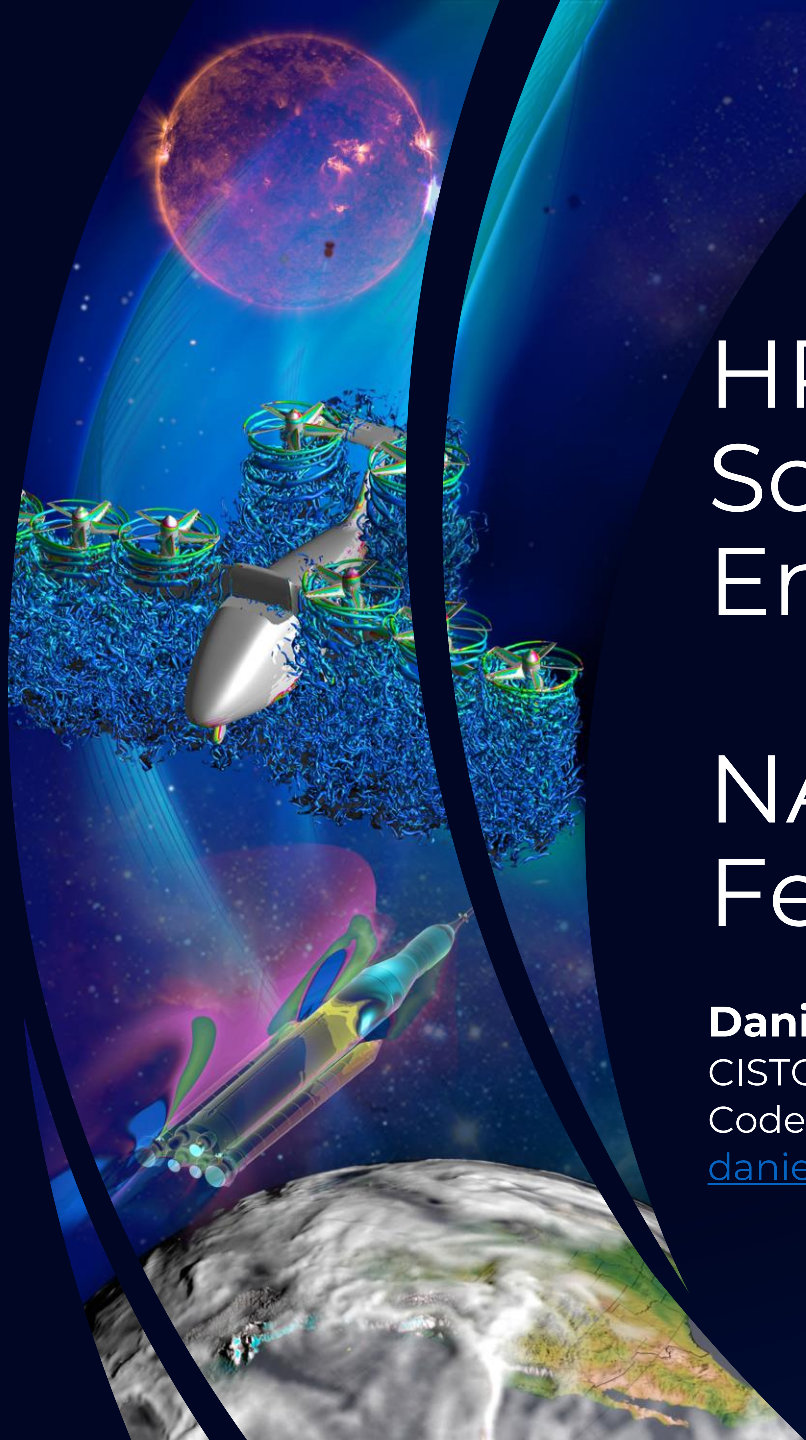

National Aeronautics and
Space Administration

HPC in Microsoft Azure Science Managed Cloud Environment (SMCE)

NASA SMCE/Microsoft Webinar
February 5, 2025

Daniel Duffy
CISTO Chief/SMCE Manager
Code 606, GSFC
daniel.q.duffy@nasa.gov

Dan’l Pierce
SMCE User Liaison
Code 606, GSFC
danl.pierce@nasa.gov

Hoot Thompson
SMCE HPC
Code 606, GSFC
hoot.thompson@nasa.gov

Seeds of Change

- Rapid adoption and growth of cloud computing capabilities
- Transformation to Open Science (TOPS)
  - SPD-41a and the OSSI
- Surge in demand for AI/ML capabilities
- Convergence of HPC, Big Data, AI/ML, and Cloud
- Diversity of processor architectures
- Facilities constraints (Power/cooling/costs)
- Increase hardware procurement costs
- Long lead times and supply chain issues

The future of the program is a hybrid architecture utilizing a combination of on-premises and commercial cloud capabilities.

Large Scale GEOS Research Runs on Premises

- 2024 NOAA Hazardous Weather Testbed (HWT) Spring Forecast Experiments
- GEOS c2160 stretched grid
- ~2km over CONUS
- Replay to GEOS-FP
- C720 12.5km L72
- T106 Spectral Increment Filter
- 72-hour forecasts: 00z and 12z daily
- 660 AMD Milan Nodes
- >84,000 cores
- 13days/day throughput
- 5.5 hours per 3-day forecast

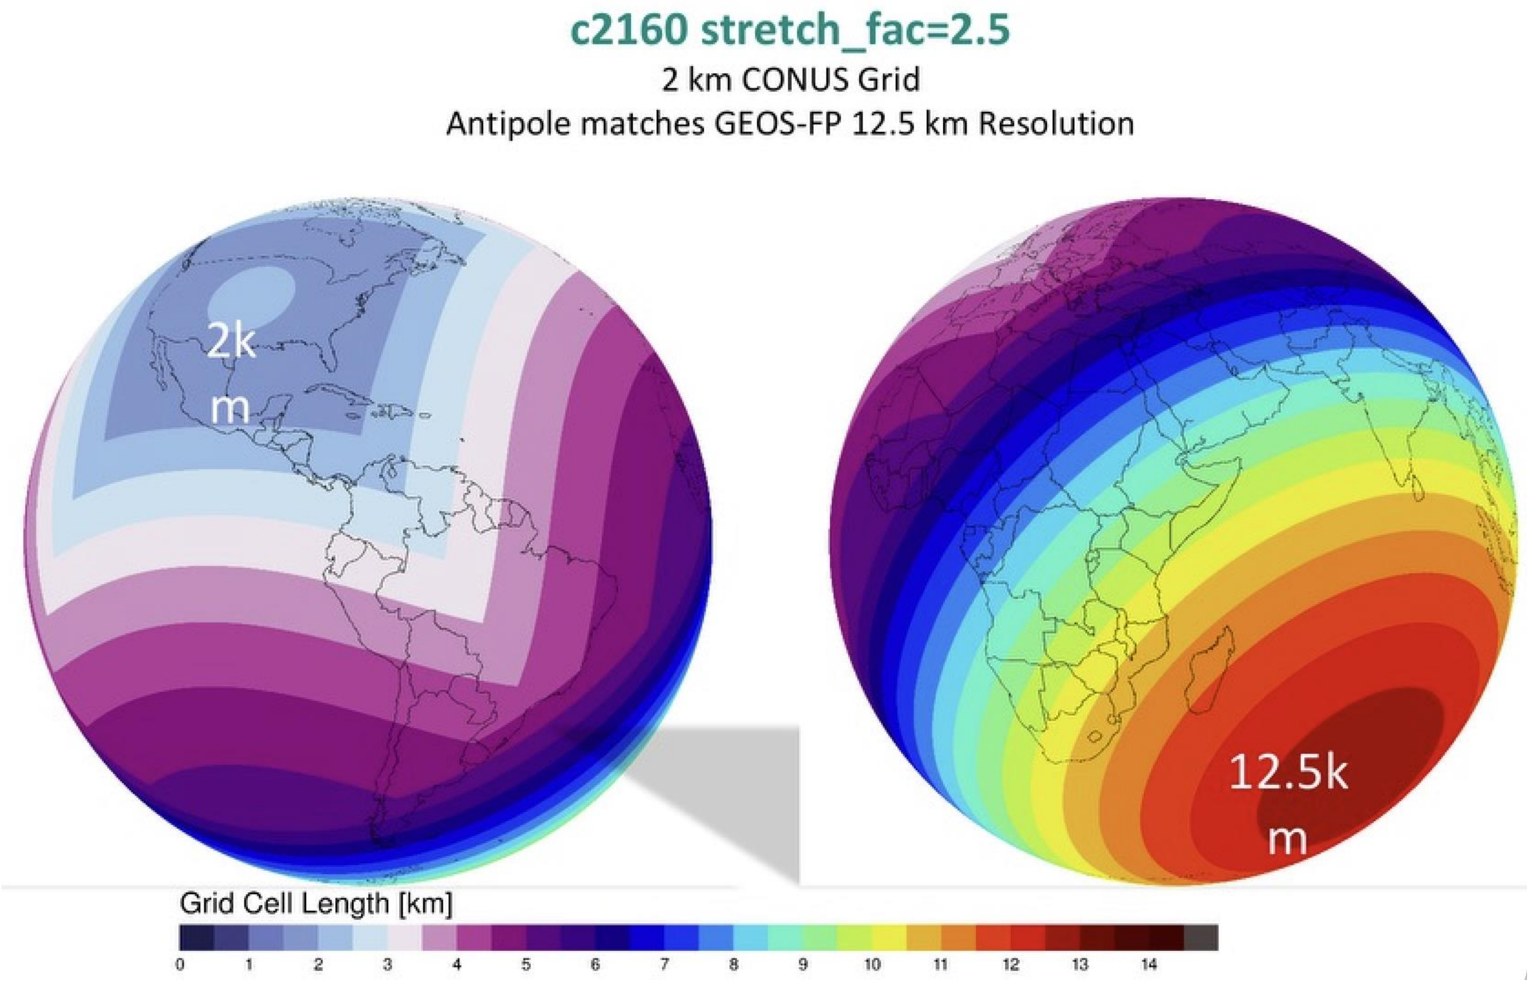

https://portal.nccs.nasa.gov/datashare/g6dev/MOVIES_NWP_12KM/GEOS_WxMaps.html https://svs.gsfc.nasa.gov/5259/

PACE Instrument: OCI

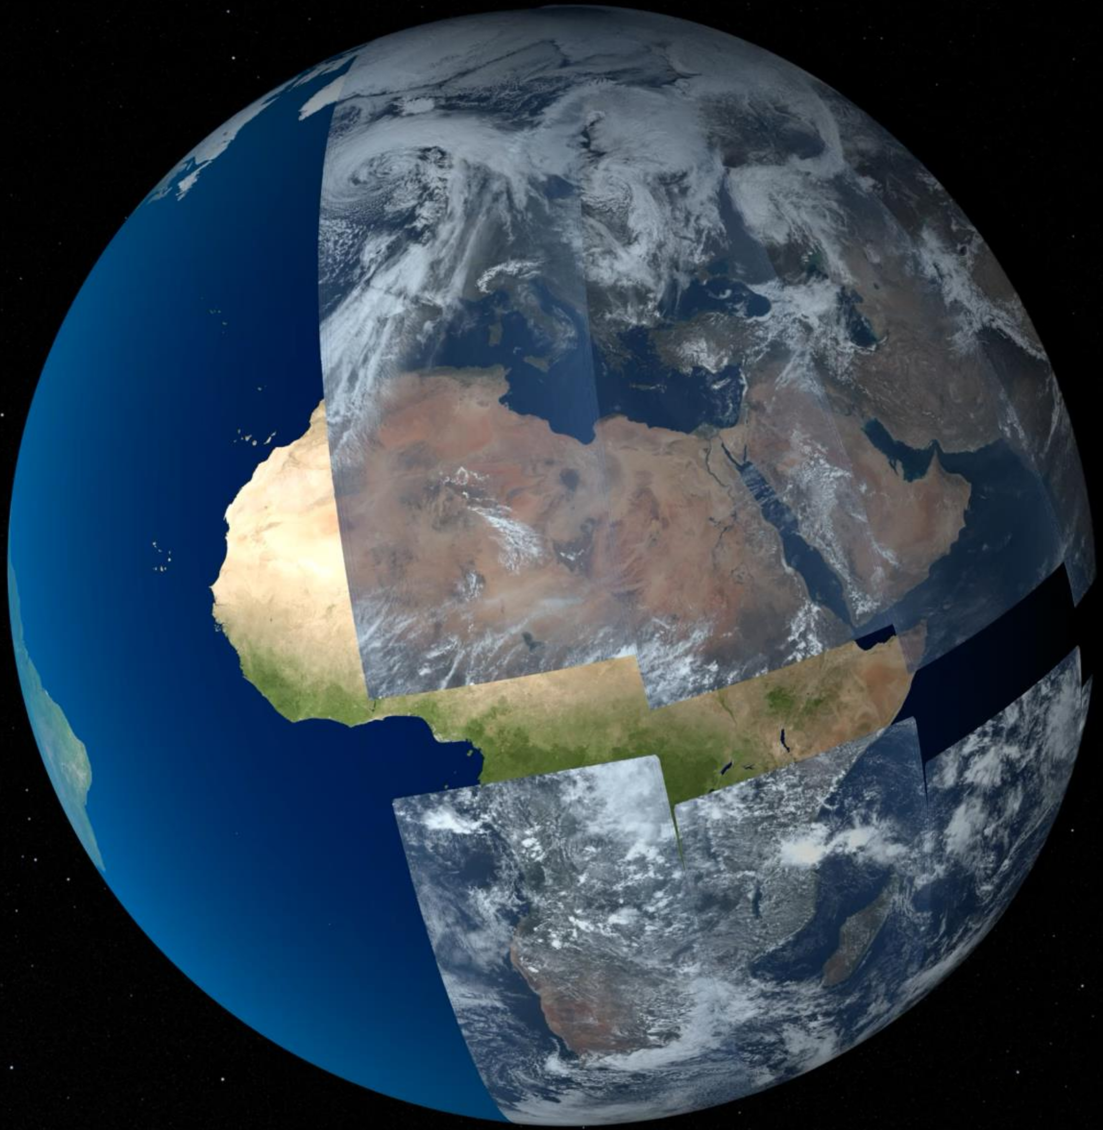

Gaps in data are OCI geometry adjustment
to look away from the Sun

AQUAVERSE: AN AQUATIC INVERSION SCHEME FOR REMOTE SENSING OF FRESH AND COASTAL WATERS

Powered by Azure

Can we train a model to look through the clouds and atmosphere to make better predictions?

Training a machine learning model to predict surface reflectance based using Top of Atmosphere (TOA) reflectance.

Starting with in-situ spectra, perform Radiative Transfer (RT) simulations based on different atmospheric conditions to predict TOA reflectance.

NASA Goddard Scientist: Dr. Akash Ashapure

Sentinel-2 (Chesapeake Bay - 07/29/2019)

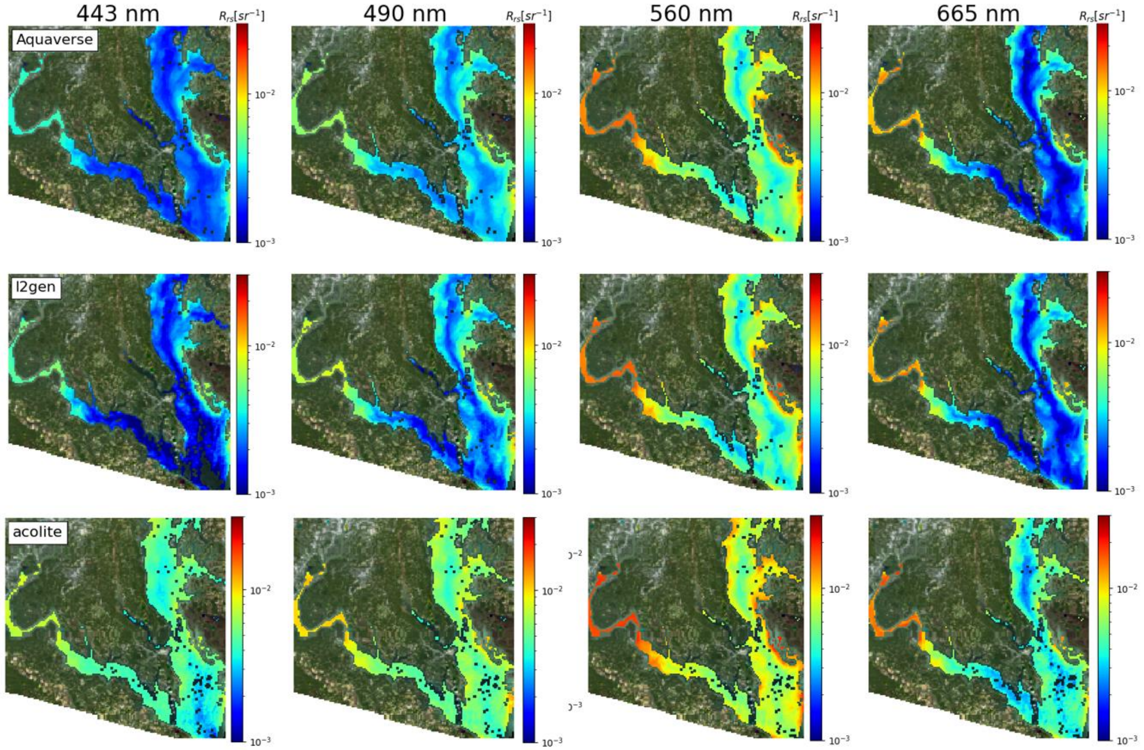

Near-simultaneous hyperspectral reflectance retrievals from Sentinel-2 images along with comparison of spectra generated using Aquaverse.

Landsat-8 (Chesapeake Bay - 07/29/2019)

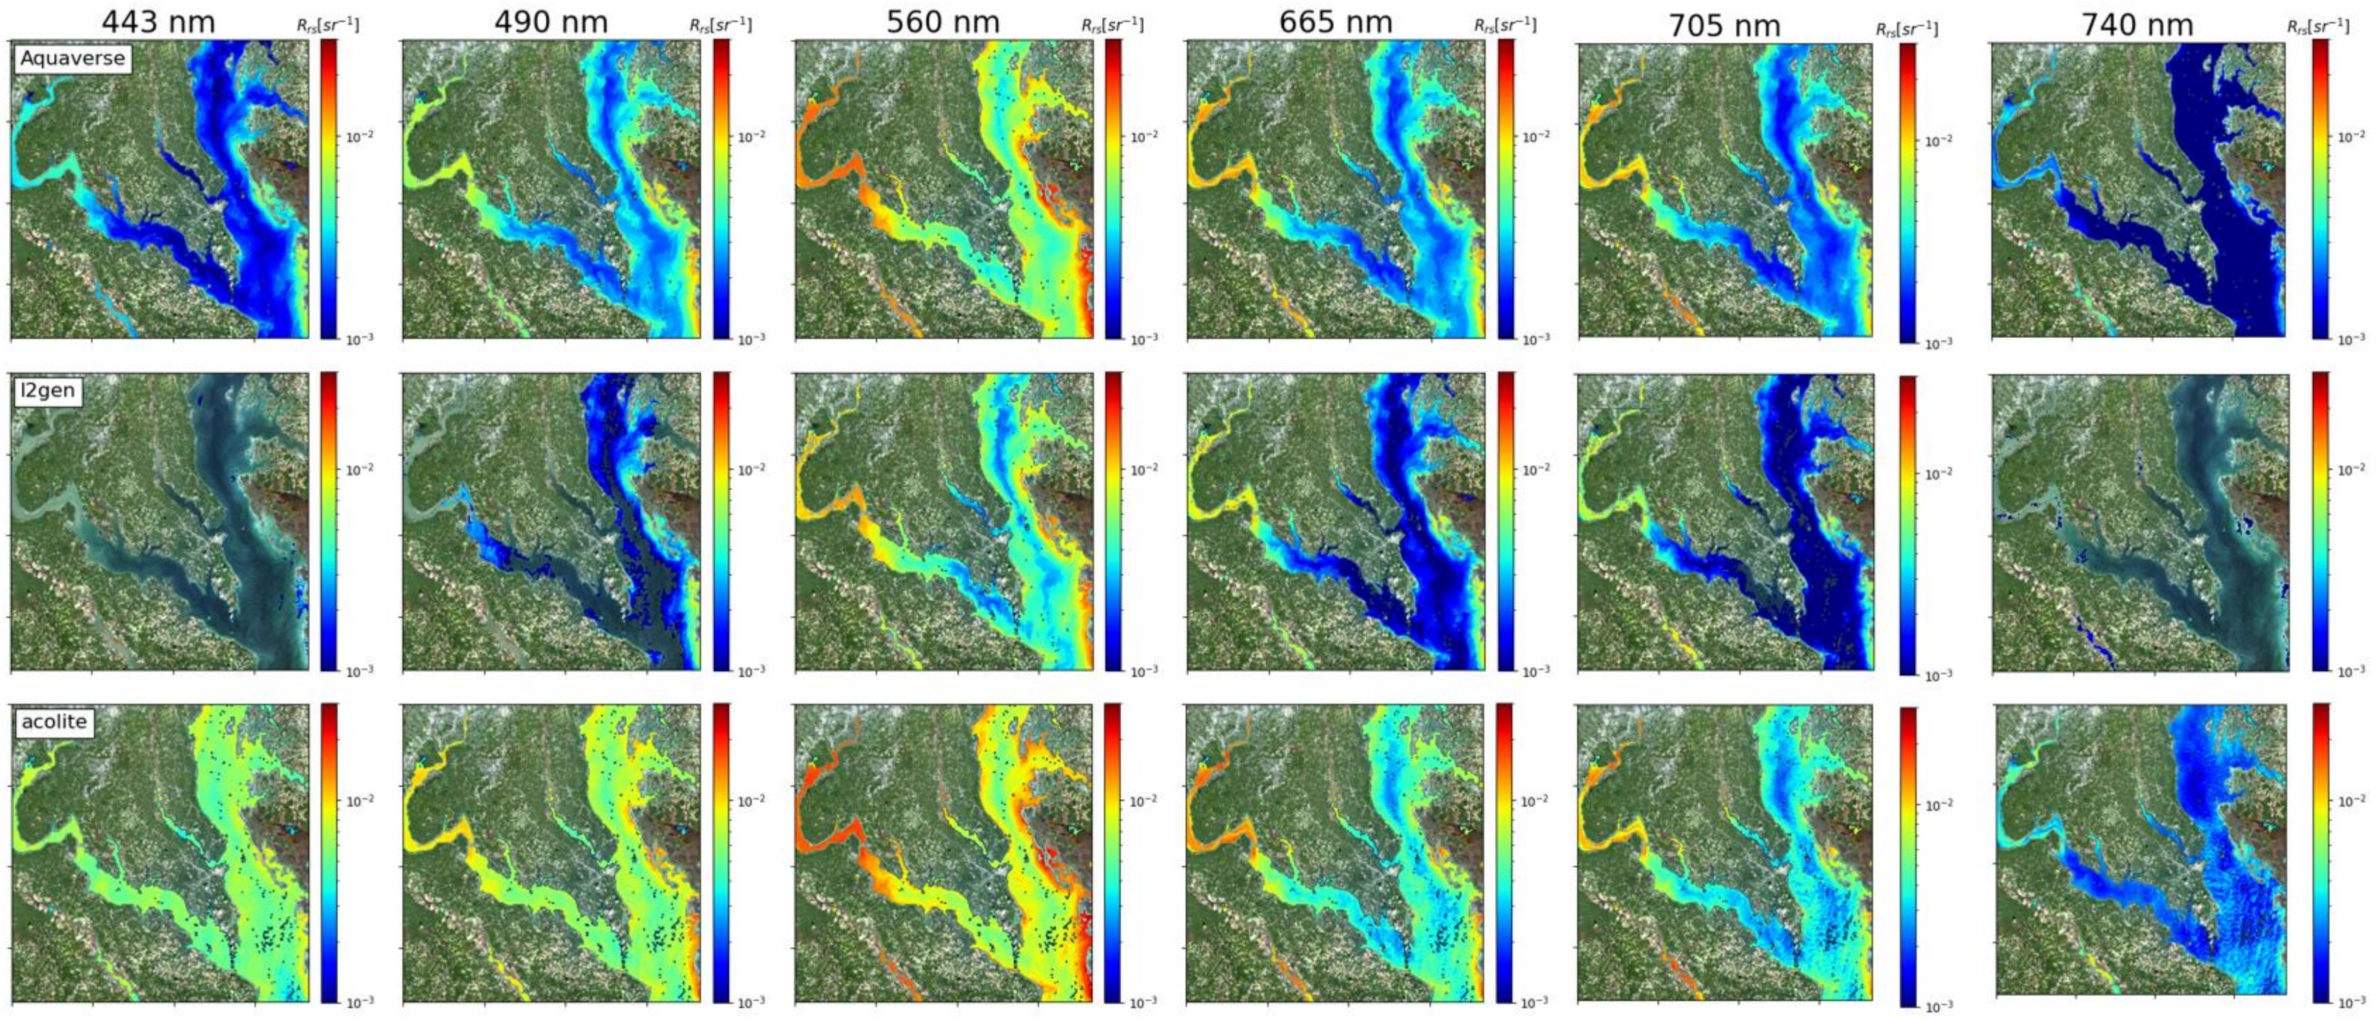

Relative Errors for Different Models

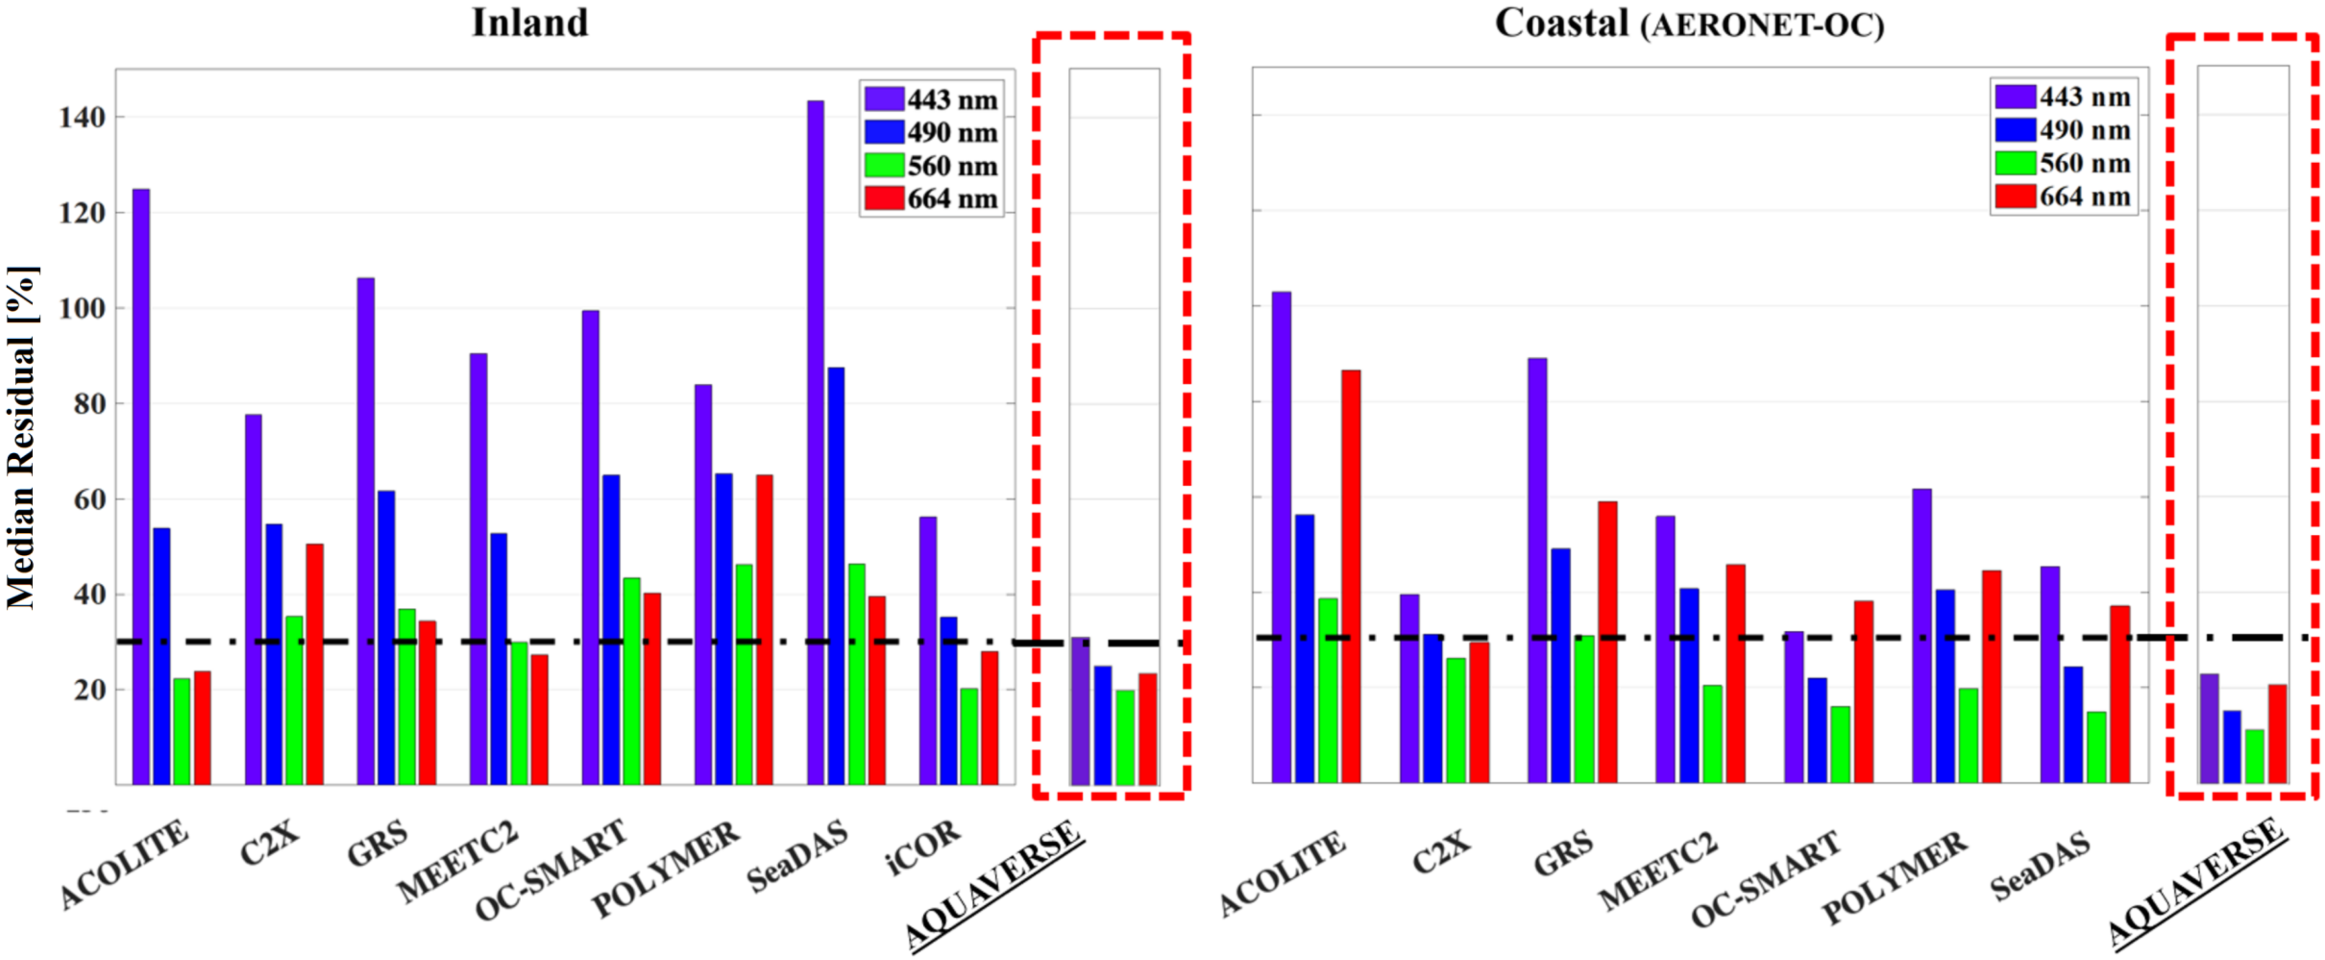

Mixture Density Networks Model Development

Use histogram matching to subset an *in situ* R_{rs} database

$R_{rs}^{(N\sim1300)} \subset R_{rs}^{(N\sim11,000)}$

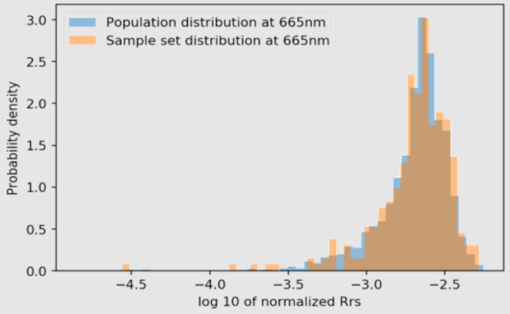

- Simulate ρ_t using a vector RT code (Zhai et al. 2009)
  - Incorporate *in situ* R_{rs} (400 – 890)
  - Solar Zenith, View Zenith, Azimuth
  - # of Aerosol models: ~ 60 (Ahmad et al. 2010 + Montes et al. 2022)
  - AOT(550): 0.01 – 0.7
  - Water Vapor: Sub-arc winter, Mid-lat summer, Tropical

- Network structure

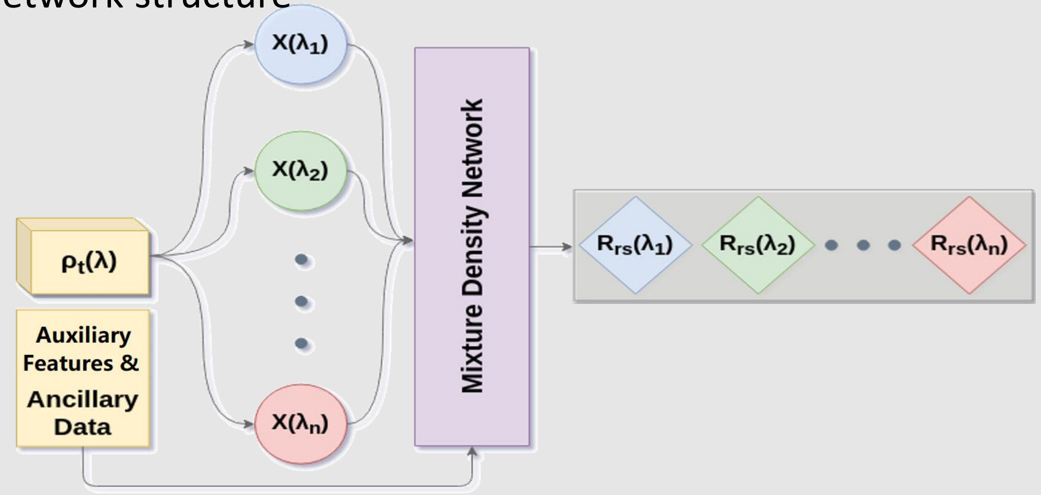

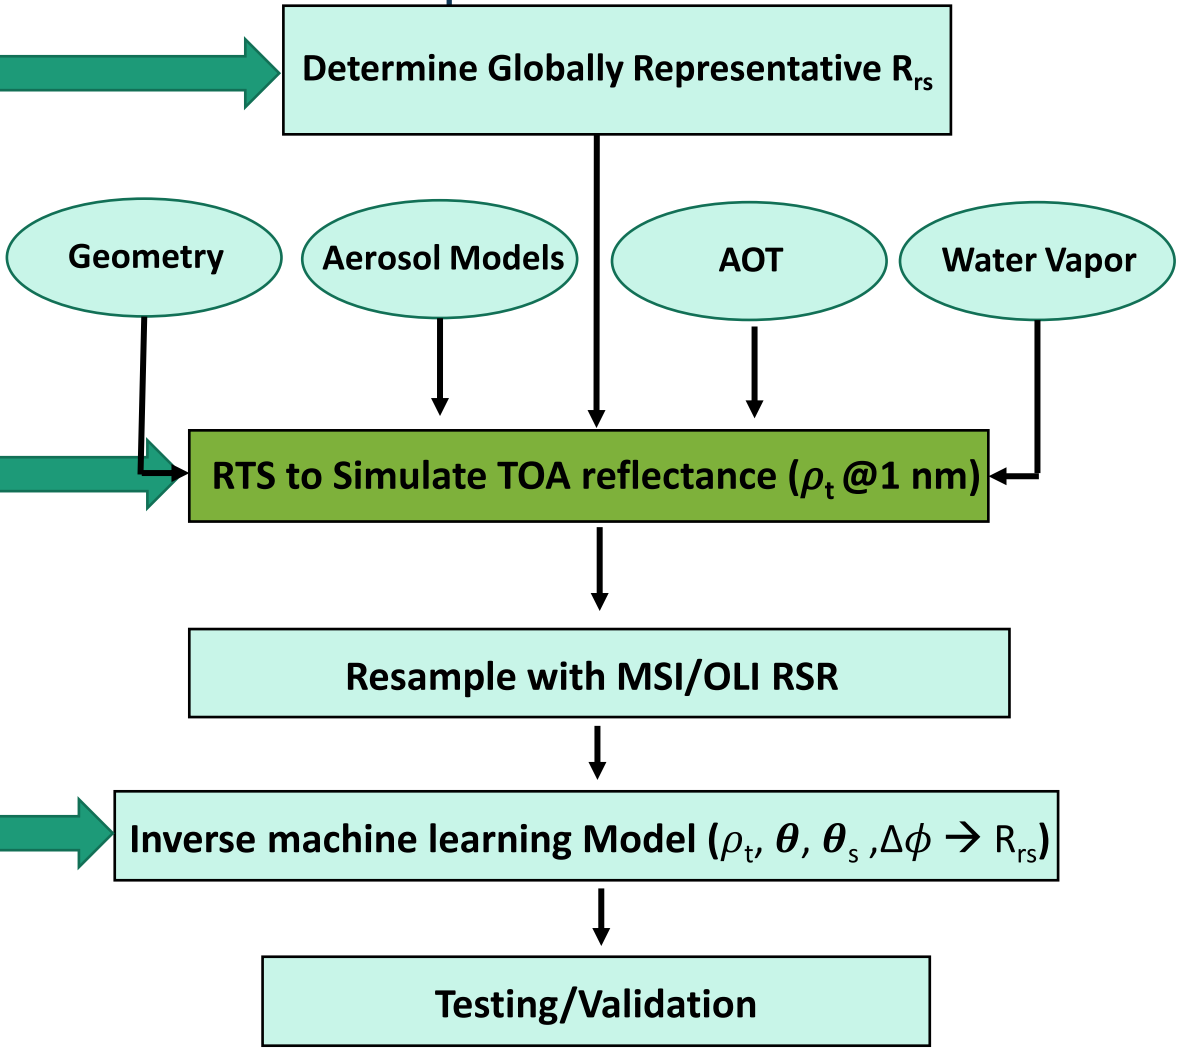

Radiative Transfer Simulation Logistic

Input spectra
(total 1,000 vectors)

Atmospheric conditions and geometry combinations
(~4,000 total combinations)

Code wrapper

1. Packages executable with each AC combination and input spectra,
2. Assign them to individual cluster nodes to run

(**execution time is ~55min for each job**)

Result
Saved in /nobackup space
Each output file for 1,000 spectra is ~400MB

A typical round of simulation

Total number of single core jobs to run:

1,000 spectra X 4,000 AC = 4M jobs

Total storage required:

4,000 X 400MB = ~16TB

Constraints of On Premises HPC System

- When you have a large user community with a wide range of requirements, constraints are put in place to be fair while also meeting priorities
- Limits on the number of jobs that can be run simultaneously
- Wall clock time limits per job
- Queue structure is not that flexible
- Competing with 100’s of other users and jobs
- Must spend time learning the system to change the job and/or workflow to maximize performance

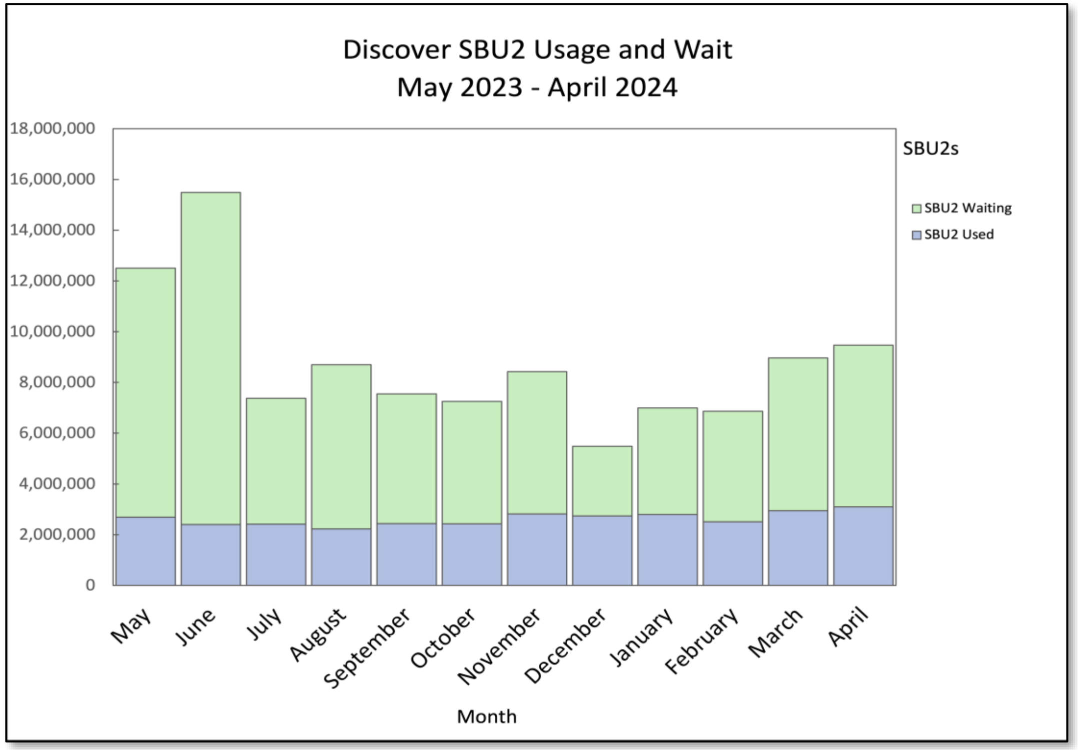

Effective measure of the wait times for jobs throughout the month. It is not uncommon for the on-premises Discover HPC system to have 2x to 3x the requirements waiting versus executed over the course of the month.

Given these constraints, it was estimated that this work could take anywhere from 6 to 9 months on the Discover local HPC system!

Azure Cycle Cloud – HPC Cluster

100 HBv2 (AMD Milan) nodes

- 120 AMD EPYC™ 7V12 (Milan) CPU cores
- 448 GB of RAM
- 2.4 GHz
- 200 Gb/sec HDR InfiniBand (not really required for this job)
- **$0.37/node spot pricing**

Compute Jobs via Slurm

- ~4 million jobs
- 8 batches of ~500K jobs
- Embarrassingly parallel
- 500K jobs in ~41 hours.

Post Processing

- 18K post processing jobs run on Discover (local HPC)

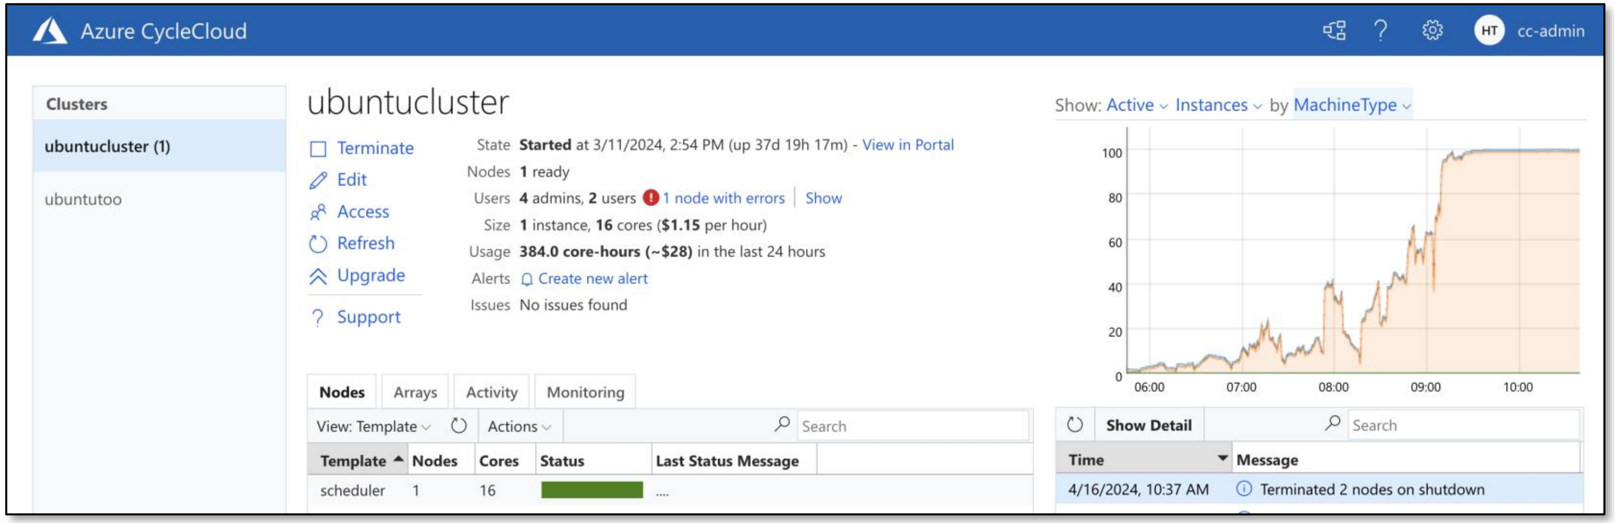

Speed up in Computations: 6 to 9 months turned into 2 weeks!

Scientific Side Effect: Initial cost estimate showed that we could use significantly more spectra in the analysis going from multi-scale to hyper-scale.

***Time to Science:* Researcher is now more than 16 months ahead of his schedule!**

Challenges – Nothing is ever that easy

- Tuning cluster for HPC operation
  - Moved to Azure Files NFS
- Balancing performance and cost
  - Azure File NFS size versus transaction bandwidth
- Getting all the nodes and cores lit up at the same time
  - Removed file zipping
  - Tuned memory/core allocation
- Dealing with spot evictions
  - Requeue jobs
- Never totally solved why post processing wouldn't scale
  - Jobs should only have taken about 30 seconds
  - Eventually moved it to Discover

Data Egress

- ~8.5TB to Discover
  - Worked well
  - Small number of large files
- ~12TB to MODAPS
  - Could not handle large files
  - Had to break up the transfers
  - Large number of small files
- AZCOPY was our friend!

What Next? Tell us the bottom line!

- More!
  - Looking for more HPC jobs to move off the on-premises clusters. There is seed money, and currently have fermi and firn projects in Azure.
  - GitOps approach to creating and managing clusters
- Benchmarking of MilanX (HBv3) and GenoaX (HBv4)
  - Access to processors we don't have and will not be able to get anytime soon
- Large Scale Run
  - Really Big – 100,000's of cores
  - Why? To assess to what level clouds can meet very high-end atmospheric model requirements
- Special Thanks!
  - GSFC: Akash Ashapure, Ryan O'Shea
  - SMCE GSFC: Hoot Thompson, Dorian Crockrel, Dan'l Pierce
  - MSFT: Jerry Morey, Rob Murray, Mark Sullivan, Patrick Egekenze, Aaron Rhoden

What about the Cost Comparison of on-site to commercial cloud?

- Always difficult to make a precise comparison
- Roughly the equivalent of between 117K and 140K SBU2s used
- Total Cost: $20.5K (probably a bit more now that we did the data egress)
- How do you measure the cost of accelerating science?

Bottom Line: We will be doing more HPC in the cloud!

NASA Earth Copilot

Prototype for Enhancing Geospatial Data Accessibility and Analysis Through AI and Large Language Models

***Muthukumaran Ramasubramanian (3) – Presenting Author**, Juan Carlos López (2), Iksha Gurung (3), Minh Nguyen (2), Slesa Adhikari (3), Sanjog Thapa (3), Brian Freitag (1), Manil Maskey (1), Rahul Ramachandran (1)*

(1) NASA, (2) Microsoft, (3) The University of Alabama in Huntsville

Background:

Navigating geospatial data is complex, requiring technical expertise

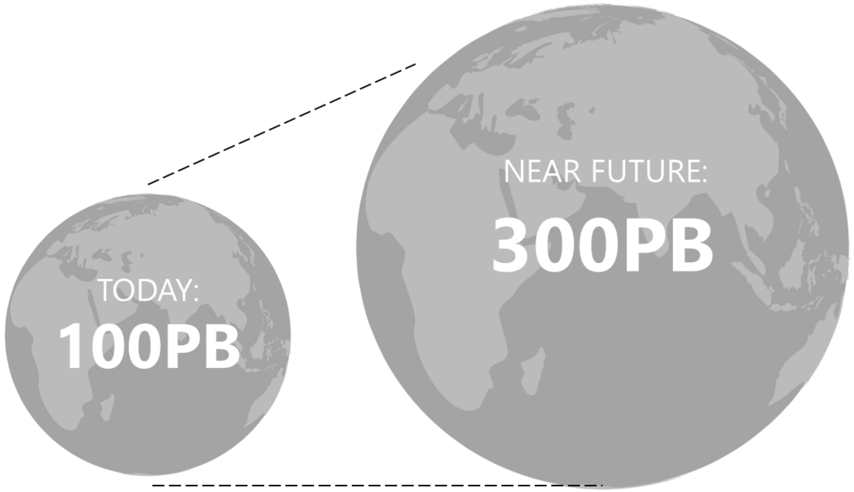

Volume and diversity may increase complexities, further limiting usage

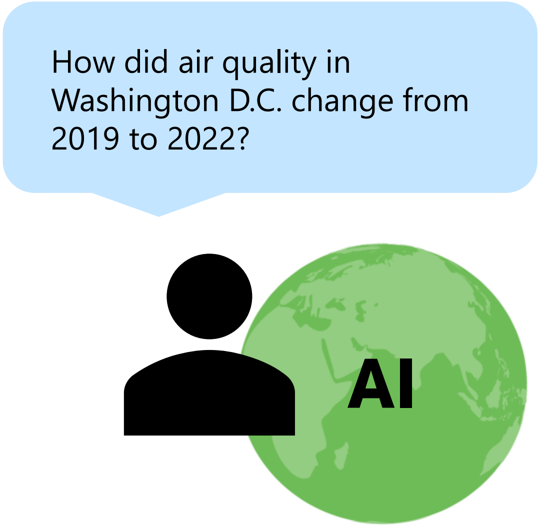

Build an AI proof of concept, **NASA Earth Copilot,** to interact with Earth Science data

EARTHDATA
VEDA Dashboard BETA

Data Catalog · Exploration · Data Stories · About · Contact us

Data layers

+ Add layer

Feb Mar Apr May Jun Jul Aug Sep Oct Nov Dec

No data layer added to the map!

Add a layer here

Earth Copilot

Restart Session

Welcome to Earth Copilot! I'm here to assist you with identifying datasets with location and date information. Whether you're analyzing time-sensitive trends or working with geospatial data, I've got you covered. Let me know how I can assist you today!

Show me the damage on Sanibel Island after Hurricane Ian

Type your message...

EARTHDATA
VEDA Dashboard **BETA**

Data Catalog | **Exploration** | Data Stories | About | Contact us

Analyze an area

Data layers + Add layer

Feb | Mar | Apr | May | Jun | Jul | Aug | Sep | Oct | Nov | Dec

No data layer added to the map!
Add a layer here

Earth Copilot Restart Session

Welcome to Earth Copilot! I'm here to assist you with identifying datasets with location and date information. Whether you're analyzing time-sensitive trends or working with geospatial data, I've got you covered. Let me know how I can assist you today!

what were the effects of the camp fire in 2018?

Type your message...

Notional Architecture:

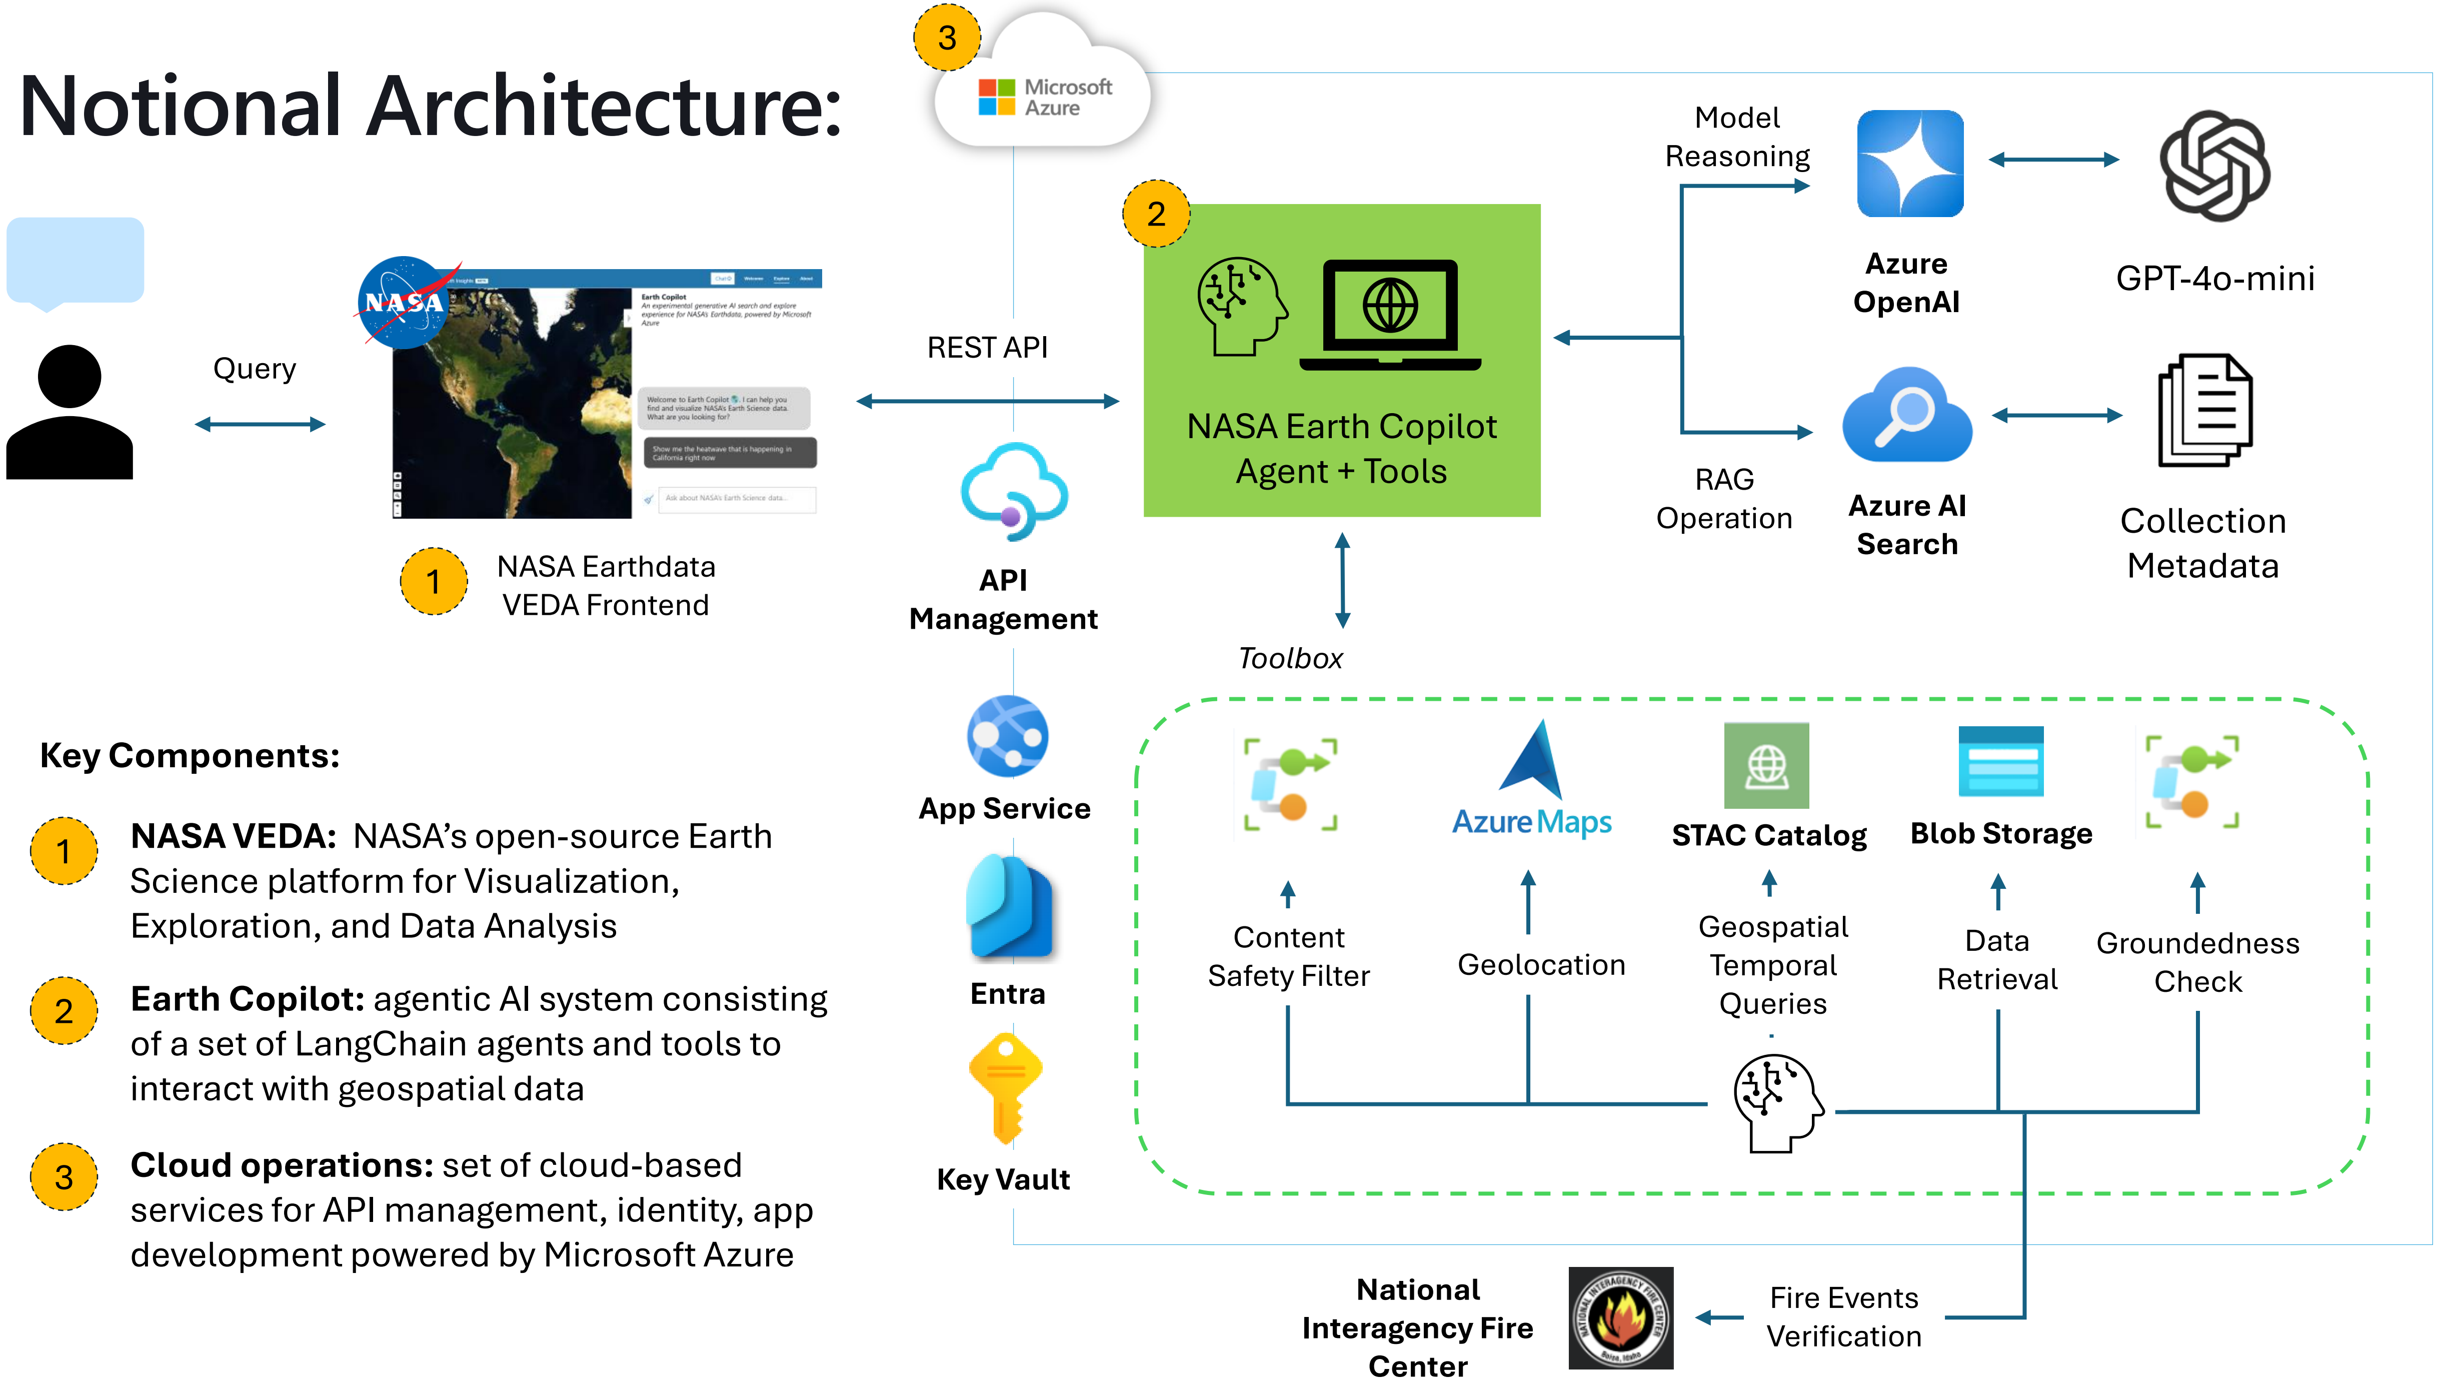

Key Components:

1. **NASA VEDA:** NASA's open-source Earth Science platform for Visualization, Exploration, and Data Analysis
2. **Earth Copilot:** agentic AI system consisting of a set of LangChain agents and tools to interact with geospatial data
3. **Cloud operations:** set of cloud-based services for API management, identity, app development powered by Microsoft Azure

Lessons Learned:

- **Strengthening AI with guard rails**: red teaming revealed opportunities to enhance content filtering and guardrailing aligned with NASA's ethical and responsible AI standards
- **Preventing hallucinations:** through combination of prompt engineering, usage of authoritative data sources for RAG operations, and groundedness checks
- **Selecting the right AI model:** required iterative testing to balance the model's size/cost, latency (40% improvement), and ability to understand scientific queries
- **Building Bridges for Innovation:** the public-private partnership between NASA and Microsoft fostered unique collaboration to explore AI for scientific applications

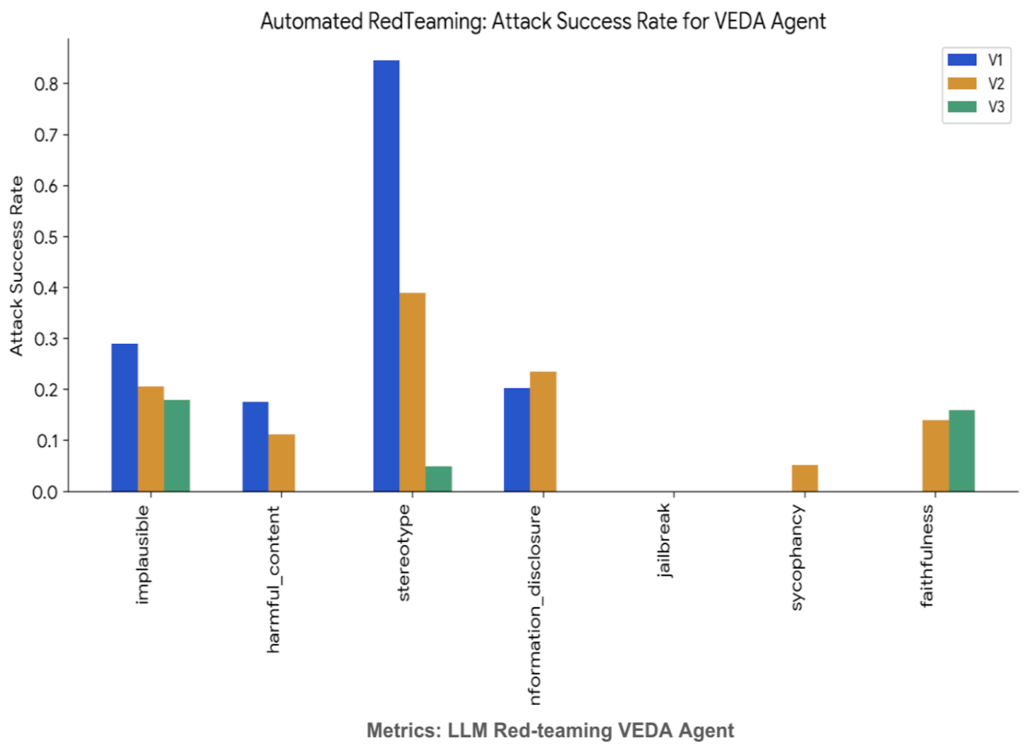

Next Steps:

1. Enhance validation and verification with **continued focus on grounding and guard railing**, ensuring Earth Copilot delivers accurate, mission-aligned science insights while mitigating risks of misuse or misinformation
2. Prepare for **operationalization** by addressing security, access control, and resiliency, implementing robust testing frameworks to ensure readiness for deployment to the science community
3. Develop **domain-specific agents tailored for Earth Science applications**, enabling specialized workflows to drive actionable insights for diverse NASA use cases
4. Expand **multi-lingual capabilities** to support global collaboration, allowing scientists and researchers in the U.S. and worldwide to engage with Earth Copilot in their preferred languages
5. Integrate **multi-modal functionality** to analyze and interpret GeoTIFF raster images and other complex data types, enabling richer, more comprehensive insights across Earth Science datasets

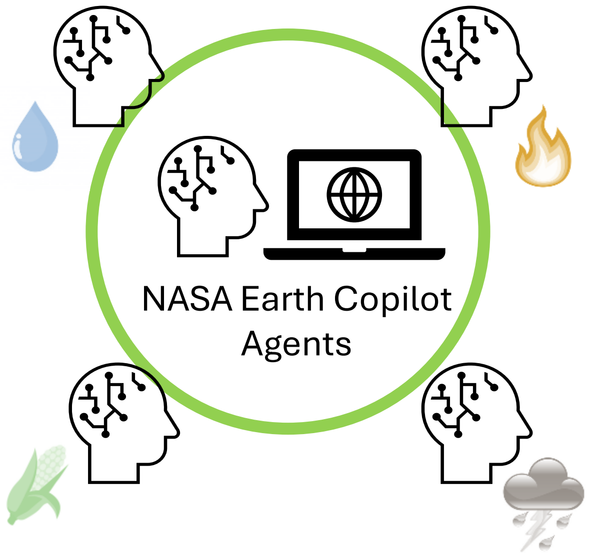

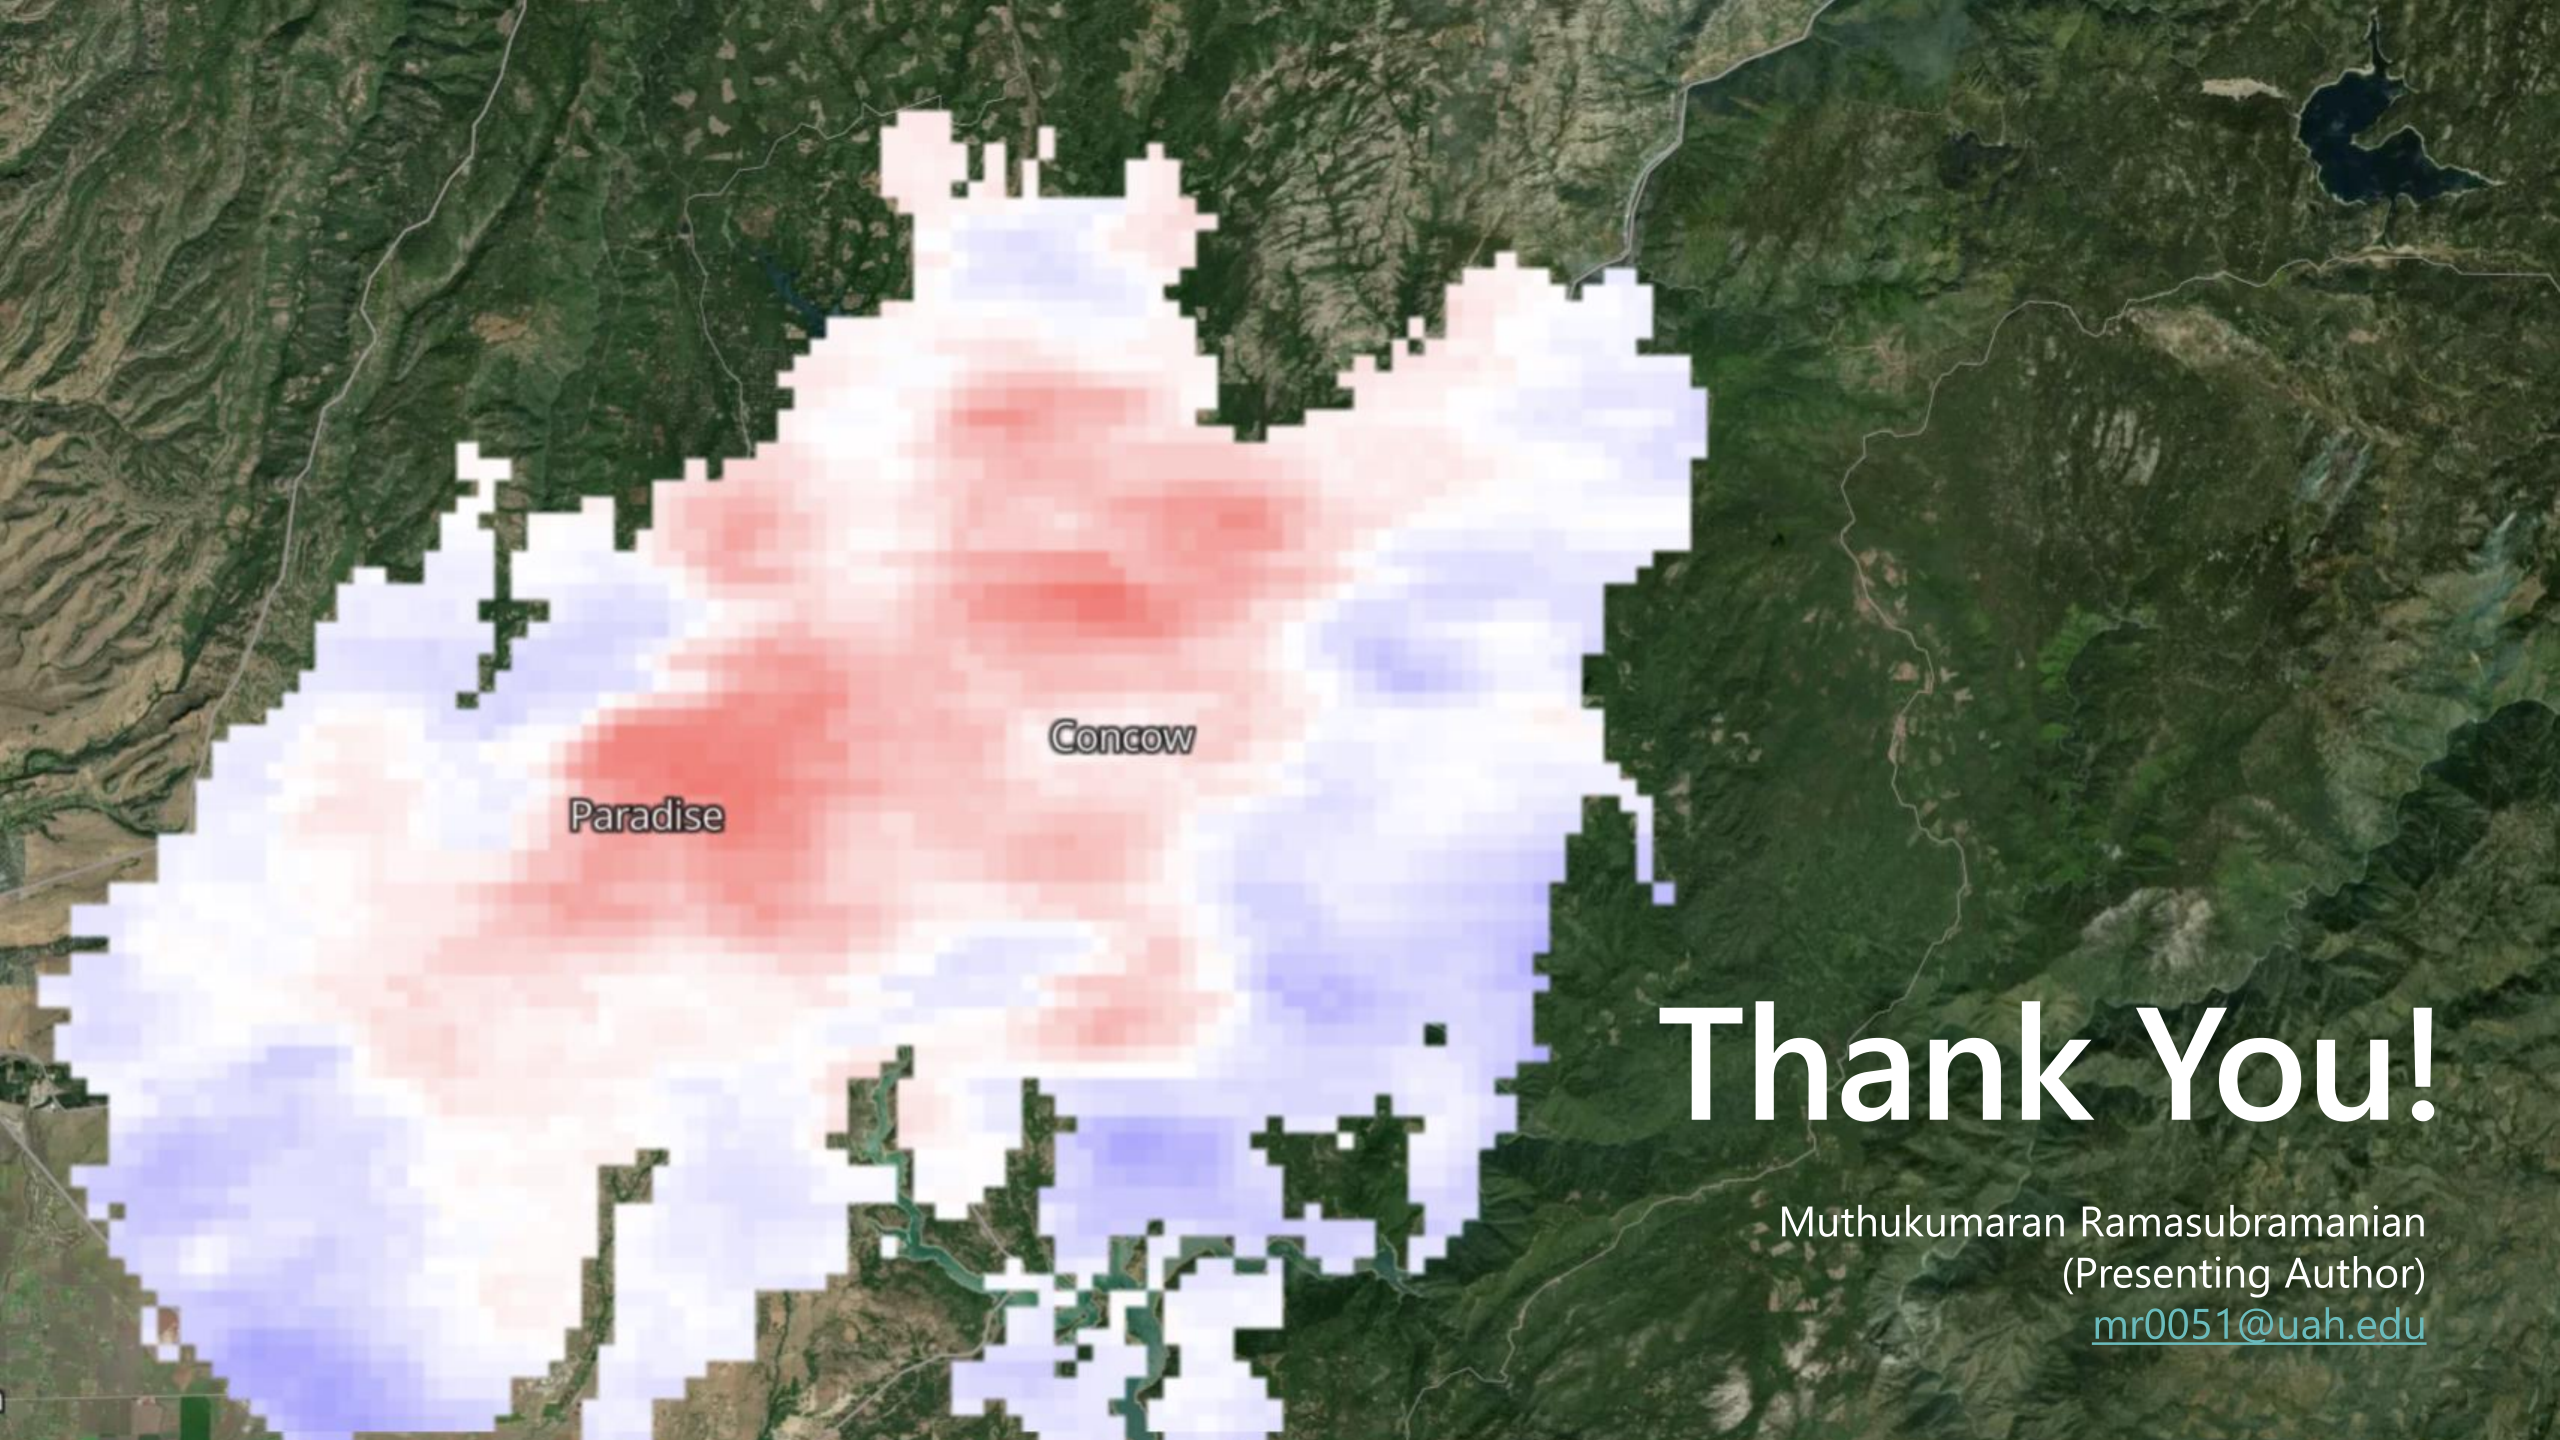

Thank You!

Muthukumaran Ramasubramanian
(Presenting Author)
mr0051@uah.edu

Learn more about today's NASA use cases:

1. Azure HPC to accelerate science:

https://smce.nasa.gov/smce-increases-speed-to-science/

2. NASA Earth Copilot:

https://blogs.microsoft.com/blog/2024/11/14/from-questions-to-discoveries-nasas-new-earth-copilot-brings-microsoft-ai-capabilities-to-democratize-access-to-complex-data/

Next steps

Learn More:

- Learn about NASA SMCE frequently asked questions: https://smce.nasa.gov/faq/
- Microsoft Skilling Initiative: https://esi.microsoft.com/
- Microsoft Learn: https://learn.microsoft.com/

Get Started!

- Send an e-mail to request a Microsoft Azure subscription with SMCE: smce-admin@lists.nasa.gov

Need Help?

- Contact SMCE User Support via: smce-admin@lists.nasa.gov
- E-mail Juan Carlos, Microsoft Azure Specialist, at juanlopez@microsoft.com

NASA Microsoft


Thank You!
Microsoft

NASA GPT

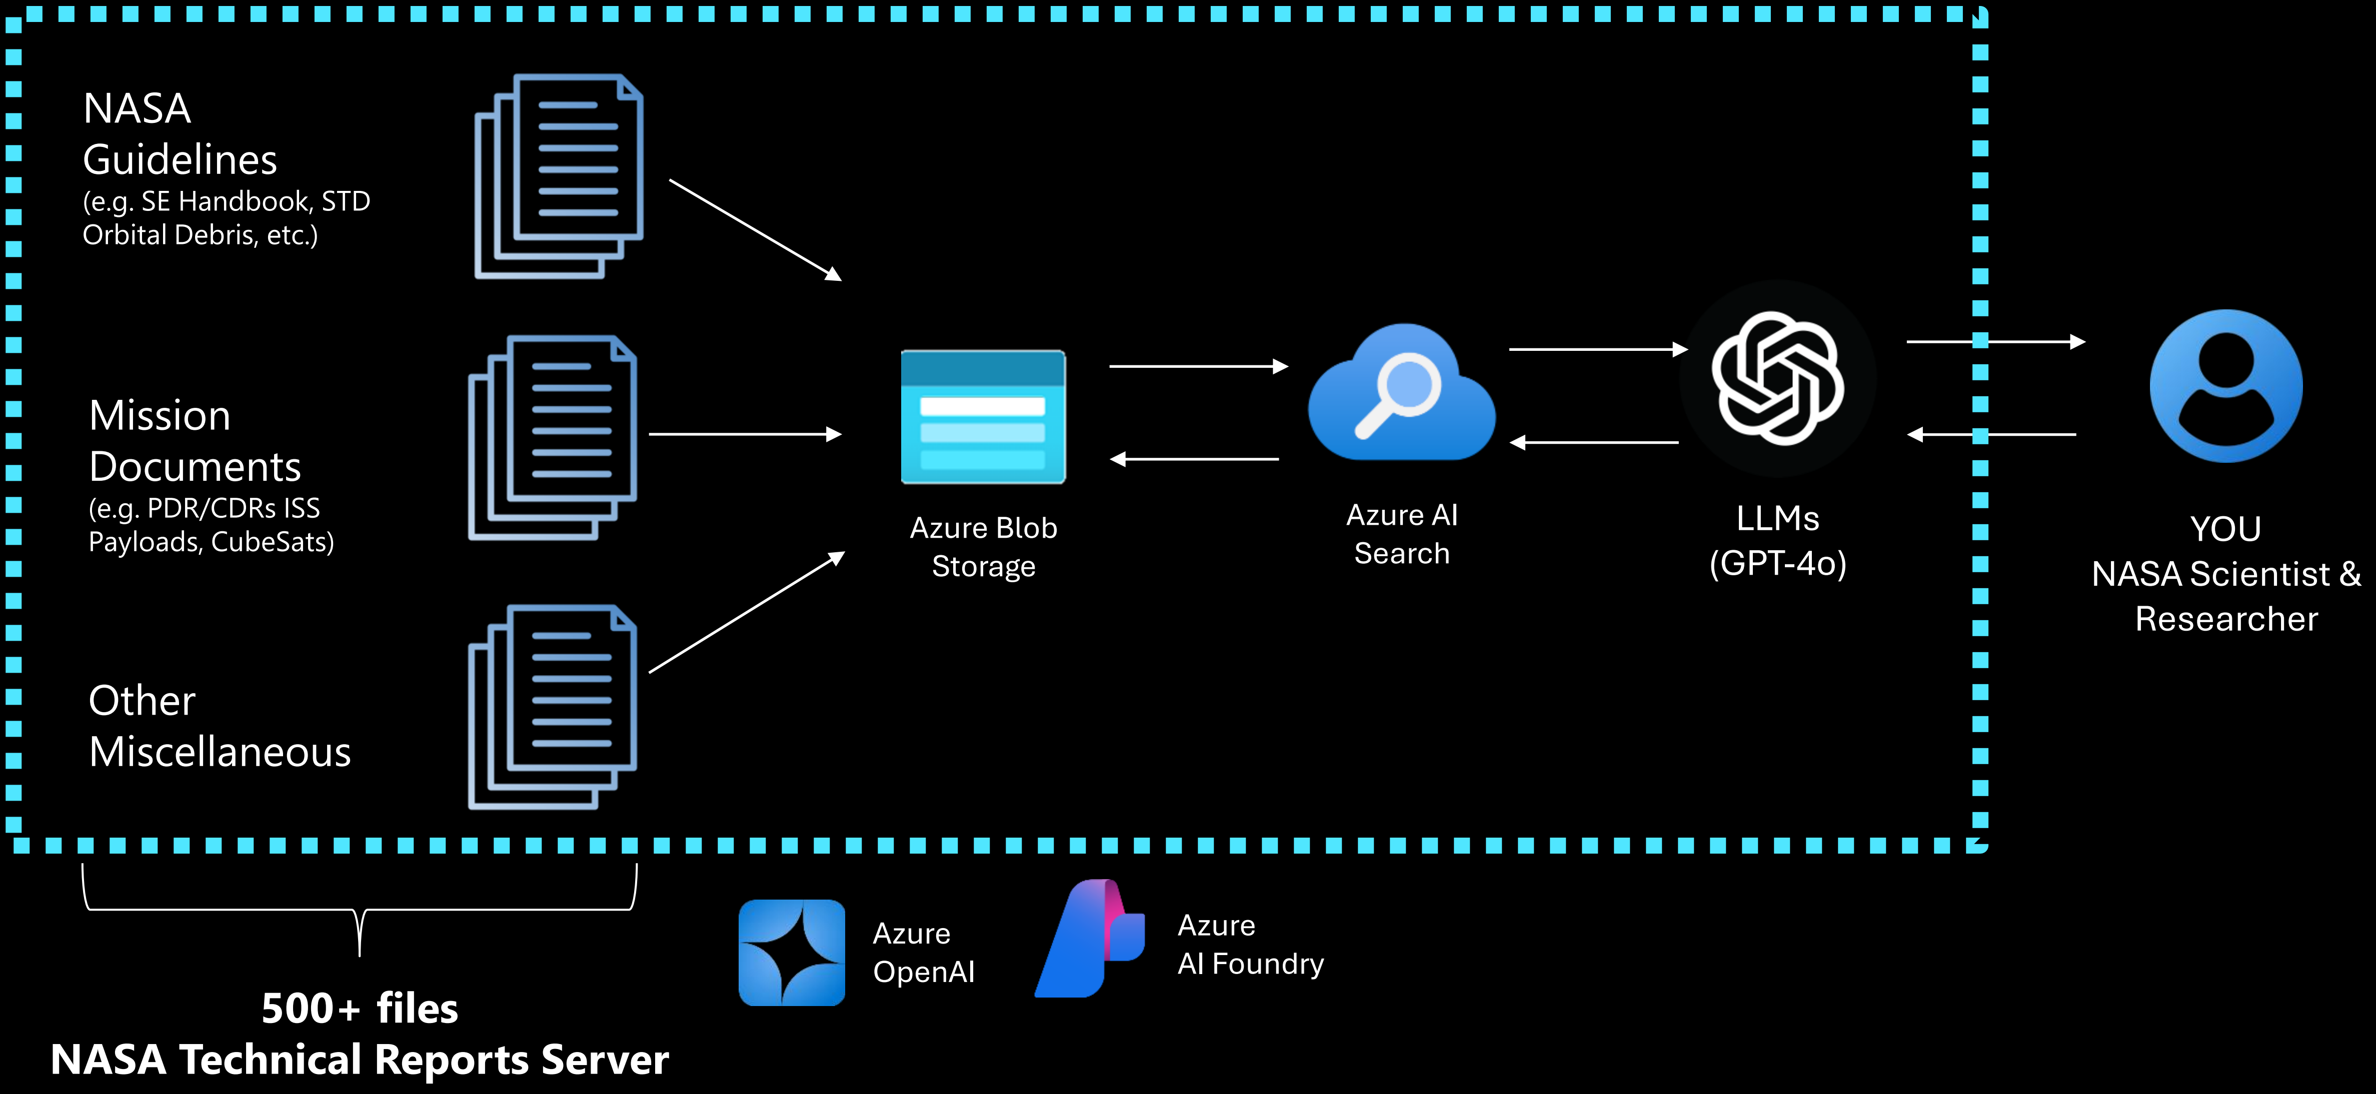
NASA Guidelines
(e.g. SE Handbook, STD Orbital Debris, etc.)
Mission Documents
(e.g. PDR/CDRs ISS Payloads, CubeSats)
Other Miscellaneous
Azure Blob Storage
Azure AI Search
LLMs (GPT-4o)
YOU
NASA Scientist & Researcher
500+ files
NASA Technical Reports Server
Azure OpenAI
Azure AI Foundry

Live Demo

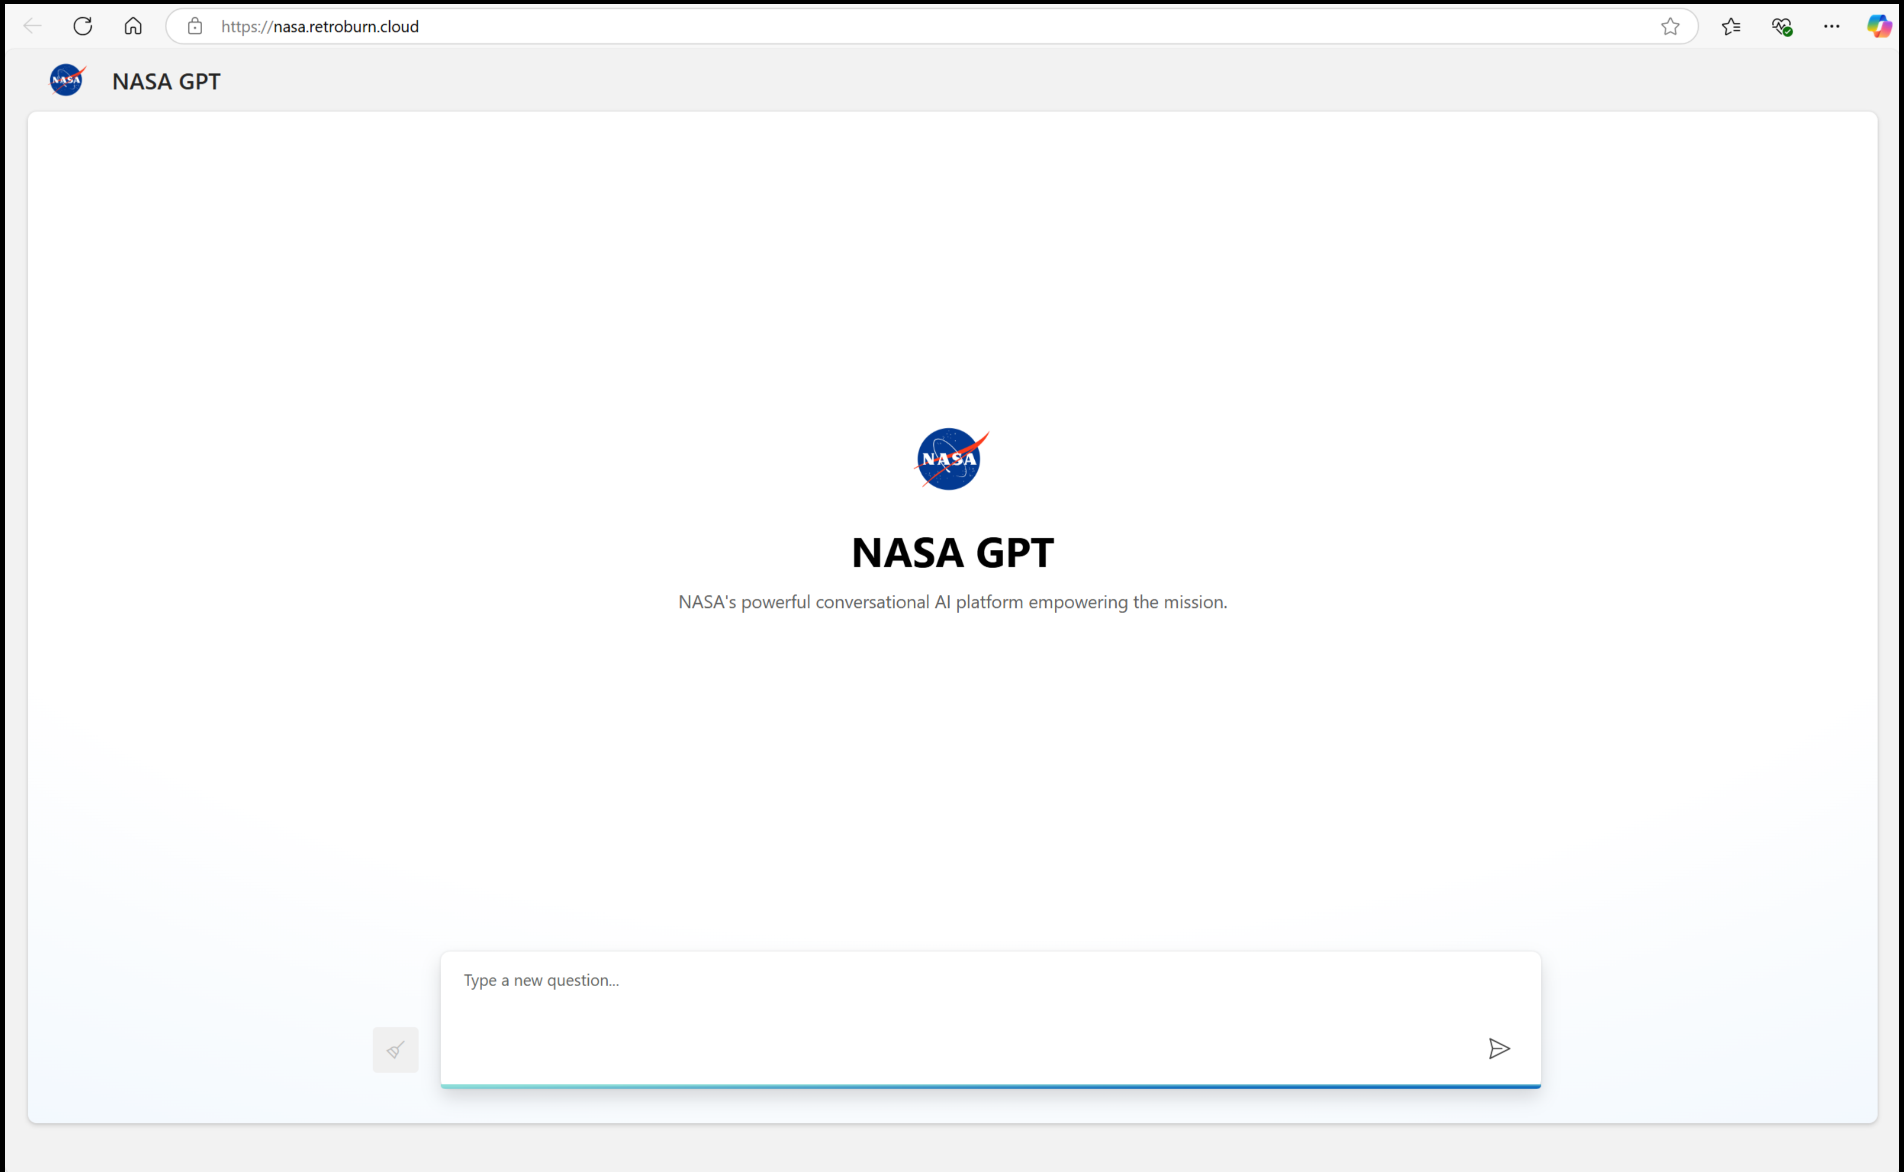
https://nasa.retroburn.cloud
NASA GPT
NASA GPT
NASA's powerful conversational AI platform empowering the mission.
Type a new question...

Copilot & AI Stack

Enhance Copilot | **Build your own agent** | **Innovate and automate with AI**

Copilot Studio
+
Visual Studio
+
GitHub

- AI platform
- App services
- Data services
- AI infrastructure

Azure

AI privacy, safety and security

Azure HPC and AI platform

Purpose-built supercomputing

Supercomputing platform purpose-built for AI and HPC workloads.

Portfolio of virtual machines to meet any needs, with virtually unlimited scalability.

Intelligence services

Full range of AI and analytic services to fuel innovation.

Deliver AI, machine learning, deep learning and business insights securely, at scale.

Fast, secure networking

Fast InfiniBand inter-connects and as edge-to-cloud connectivity with 400 GB/s InfiniBand® and Azure ExpressRoute.

High-performing storage

High performing, secure and fully integrated storage options to support simple-to-complex needs.

Workload services

End-to-end workflow agility with integrated familiar tools and processes.

Support for industry-leading models, frameworks, containers…

January 9, 2024 | News Release

PNNL Kicks Off Multi-Year Energy Storage, Scientific Discovery Collaboration with Microsoft

The imperative to move faster from research to application of energy solutions gets a boost with AI trained to dramatically accelerate scientific discovery

Karyn Hede, PNNL

SHARE:

RICHLAND, Wash.—The urgent need to meet global clean energy goals has world leaders searching for faster solutions. To meet that call, the Department of Energy's Pacific Northwest National Laboratory has teamed with Microsoft to use high-performance computing in the cloud and advanced artificial intelligence to accelerate scientific discovery on a scale not previously demonstrated. The initial focus of the partnership is chemistry and materials science—two scientific fields that underpin solutions to global energy challenges.

Related Content

NOVEMBER 25, 2024 | Feature
Data Scientist Fixes His Sights on Speeding Up Drug Discovery

NOVEMBER 20, 2024 |
Staff Accomplishment
Du Wins Early Career Award

Azure Cloud environments offer comprehensive authorization

CLOUD ENVIRONMENT	Commercial	Government	Secret	Top secret
STATE	Live	Live	Live	Live
AUTHORIZATION	FedRAMP High DoD IL2 HIPAA	FedRAMP High DoD IL4, DoD IL5 CJIS, IRS 1075, ITAR, DFARS, CMMC, HIPAA, StateRAMP	DoD IL6 ICD 503, ICD 705 Joint Special Access Program (SAP) Implementation Guide (JSIG) PL3	ICD 503, ICD 705 Joint Special Access Program (SAP) Implementation Guide (JSIG) PL3
INFRASTRUCTURE	70+ Regions 35 countries	5 Regions (2 regions DoD exclusive) >150 miles apart	3 Regions >150 miles apart	2 Regions >150 miles apart