Meeting the Big Data Challenges of Climate Science through Cloud-Enabled Climate Analytics-as-a-Service

MERRA Analytic Services

John Schnase
Office of Computational and Information Sciences and Technology
NASA Godard Space Flight Center

High-Performance Science Cloud

Dan Duffy
NASA Center for Climate Simulation
NASA Godard Space Flight Center

NASA Center for Climate Simulation (NCCS)

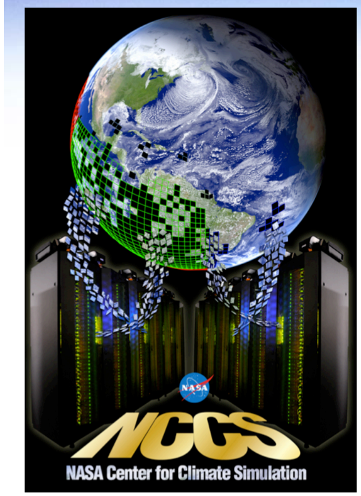

Funded by the Science Mission Directorate
- Located in B28 at GSFC

Integrated high-end computing environment designed to support the specialized requirements of Climate and Weather modeling.

Archive for Model Data ~35 PB

Assume a song is 4MB and takes 4 minutes to play

Can store 57,000 years of songs!

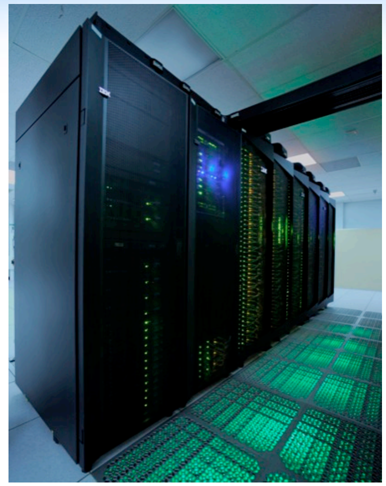

Hyperwall for Visualization

Over 15 Million pixels!

No Superbowl at this point!

High Performance Computing
- ~1.2 Petaflops
- 76th fastest computer in the world, and the entire center would be in the top 50
- www.top500.org

In one second, this system is equivalent to having every person on the earth multiple two numbers every second for over 50 hours straight!

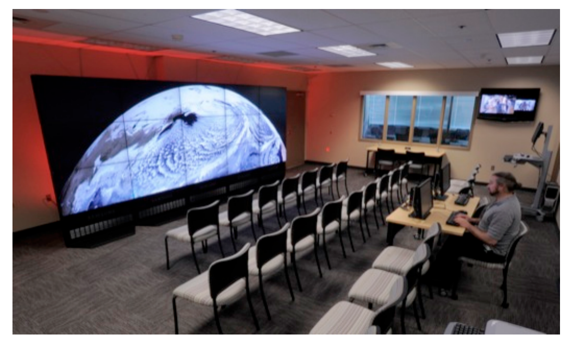

Typical HPC Applications

Takes in small input and creates large output

- Using relatively small amount of observation data, models are run to generate forecasts
- Fortran, Message Passing Interface (MPI), large shared parallel file systems
- Rigid environment – users adhere to the HPC systems

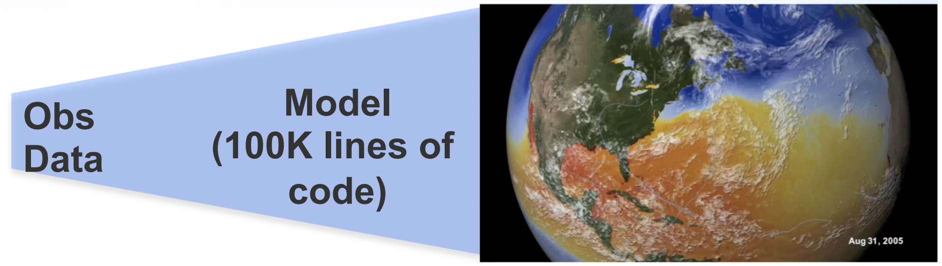

10-km GEOS-5 meso-scale simulation for Observing System Simulation Experiments(OSSEs)

Example: GEOS-5 Nature Run (GMAO)

- 2-year Nature Run at 7.5 KM resolution
- 3-month Nature Run at 3.5 KM resolution
- Will generate about 4 PB of data (compressed)
- To be used for Observing System Simulation Experiments (OSSE's)
- All data to be publically accessible
  - **ftp://G5NR@dataportal.nccs.nasa.gov/**

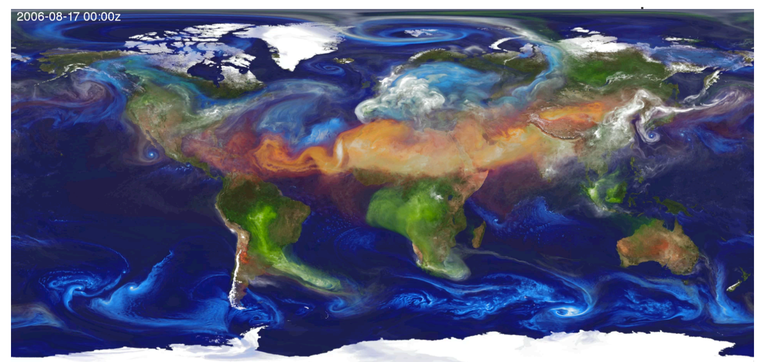

The Goddard Chemistry Aerosol Radiation and Transport (GOCART) model, Courtesy of Dr. Bill Putman, Global Modeling and Assimilation Office (GMAO), NASA Goddard Space Flight Center.

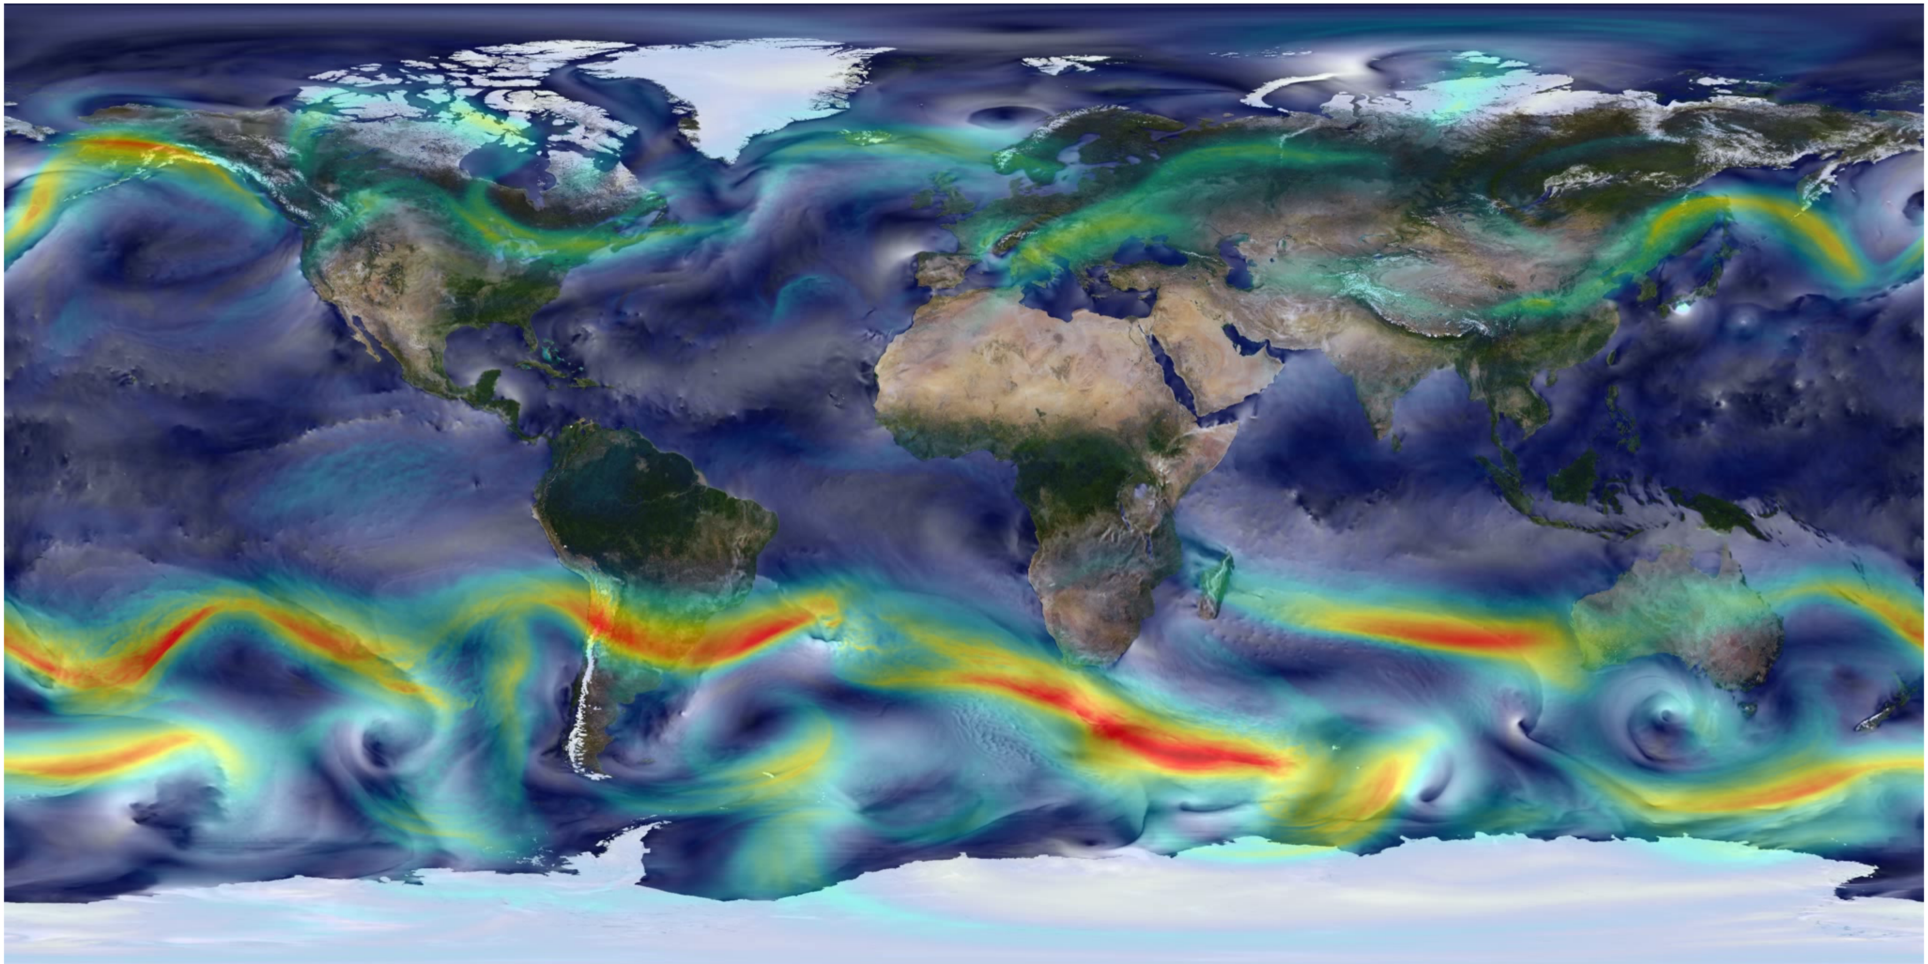

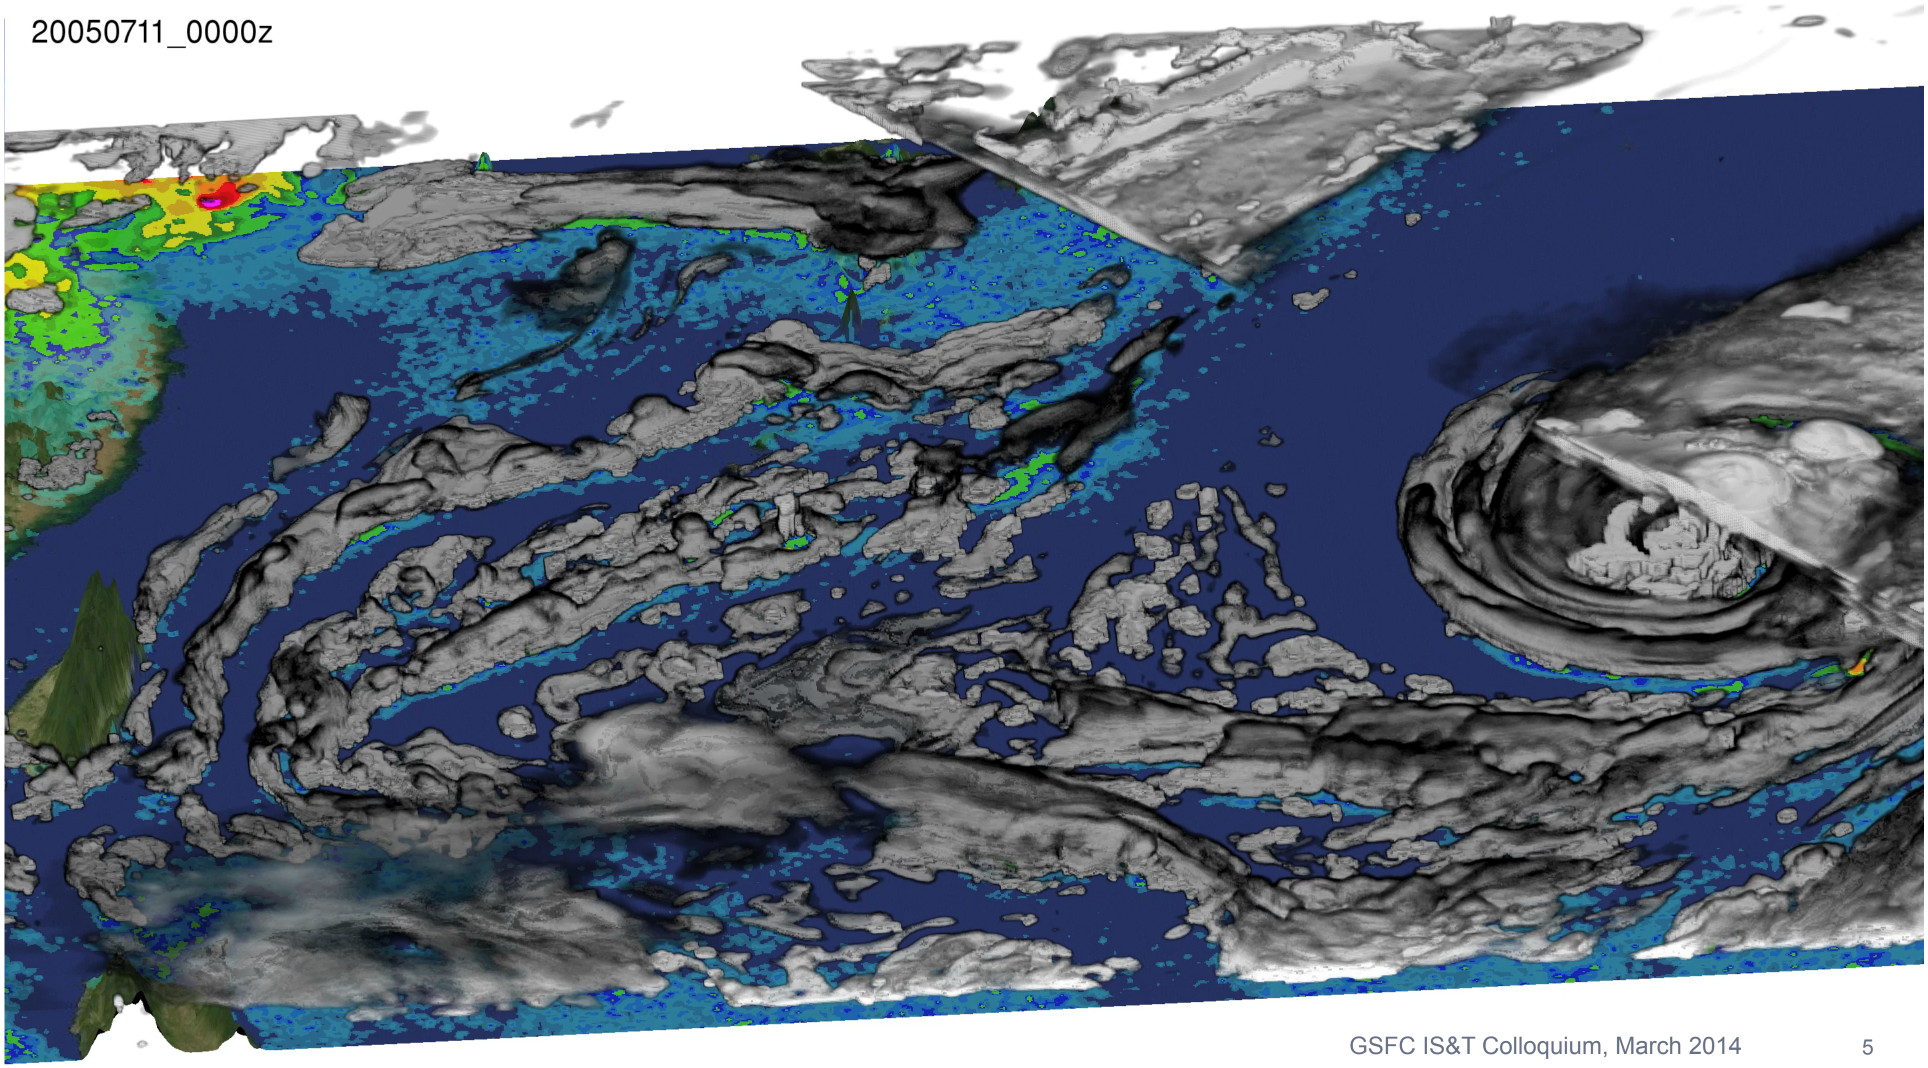
20050711_0000z

Typical Analysis Applications

Takes in large amounts of input and creates a small amount of output

- Using large amounts of distributed observation and model data to generate science
- Python, IDL, Matlab
- Agile environment – users run in their own environments

Examples

- Evaporative transport (Wei experiment)
  - Requires monthly reanalysis data sets for four different spatial extents
- Decadal water predictions for the high northern latitudes for the past three decades
  - Requires 100,000+ Landsat images and about 20 TB of storage

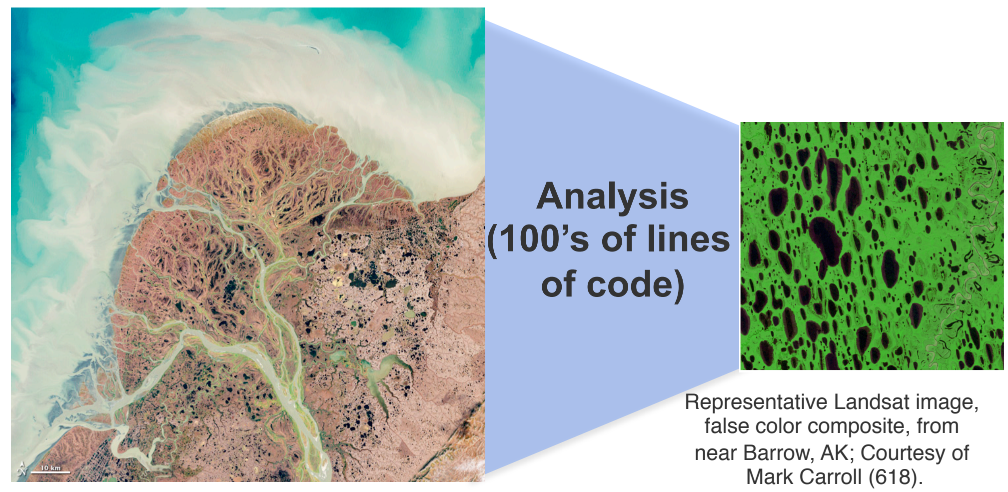

Yukon Delta Alaska; courtesy of Landsat
http://landsat.visibleearth.nasa.gov/view.php?id=72762

Representative Landsat image, false color composite, from near Barrow, AK; Courtesy of Mark Carroll (618).

Planning Science/Proposal Writing

What question am I trying to answer?

- Example: Suppose we want to generate maps of surface water from 1990 to 2012 in the arctic boreal region (problem courtesy of Mark Carroll, Code 618)

What data are available and where are they?

- Landsat time series available at the LP DAAC

How much data is needed?

- Full time series requires >100,000 scenes and ~20TB of data storage

Can I store all that data? If not, how can I process it?

- No. So download chunks of 5TB to local machine. Process. Delete. Download more.
- Projected time – 9 months – without any mistakes!

That's too long, so how can I modify my science question accordingly?

- Average across three epochs (1990, 2000, 2010)
- 25,000 scenes and ~7TB of data
- Projected time – 2 to 3 months

Scientists are limiting their questions (and science) based on the IT resources of their desktops!

Conversations Between Scientists or Conversations "We Don't Want" Between Scientists

Scientist 1

Scientist 2

Scientist 1: Hey, what are you working on these days?

Scientist 2: Oh, you know, just processing data from the new satellite for my ROSES project.

Scientist 1: You know, I need that same data for my project. Where did you get that?

Scientist 2: I downloaded it from the web.

Scientist 1: How long did it take you?

Scientist 2: Quite awhile; several weeks.

Scientist 1: Oh, man, I don't want to have to download all That data and take several weeks. Do you think I could get a copy from you?

Scientist 2: Sure, I am just not sure how to get it to you. I could NFS serve it from my machine to yours or just give you access to my system.

Scientist 1: That would be great. You don't think the security guys would mind do you?

Scientist 2: No, I'm sure they wouldn't. It is in the name of science after all.

Where do we look for help?

Google

- By 2012, Gmail had 425 million active users[1]
- Each user gets 15 GB of storage for free
- 425,000,000 * 15 GB = 6,375,000,000 GB = 6,385,000 TB = 6,375 PB = 6.375 EB
- Assuming about 6% of the email is spam[2], Gmail carried around 382.5 PB of spam!

Facebook

- By 2012, Facebook was processing 500 TB of data per day[3]
- 2.7 billion Like actions and 300 million photos per day; Facebook scanned about 105 TB every 30 minutes[4]

1. http://venturebeat.com/2012/06/28/gmail-hotmail-yahoo-email-users/
2. http://krebsonsecurity.com/2013/01/spam-volumes-past-present-global-local/
3. http://news.cnet.com/8301-1023_3-57498531-93/facebook-processes-more-than-500-tb-of-data-daily/
4. http://techcrunch.com/2012/08/22/how-big-is-facebooks-data-2-5-billion-pieces-of-content-and-500-terabytes-ingested-every-day/

Compare HPC to Large Scale Internet

High Performance Computing

- Shared everything environment
- Very fast networks; tightly coupled systems
- Cannot lose data (POSIX)
- Big data (100 PBs)
- Bring the data to the application
- Large scale applications (up to 100K cores)
- Applications cannot survive HW/SW failures
- Commodity and non-commodity components; high availability is costly; premium cost for storage

Object Storage
MapReduce
Hadoop
Cloud
Open Stack
Virtualization

Large Scale Internet

- Examples: Google, Yahoo, Amazon, Facebook, Twitter
- Shared nothing environment
- Slow networks
- Data is itinerant and constantly changing (RESTful)
- Huge data (Exabytes)
- Bring the application to the data
- Very large scale applications (beyond 100Ks)
- Applications assume HW/SW failures
- Commodity components; low cost storage

High Performance Science Cloud

Adjunct to the NCCS HPC environment

- Lower barrier to entry for scientists
- Customized run-time environments
- Code validation against older system images

Expanded customer base

- Host system for NCCS Data Services
- Science Data Processing (e.g. ABoVE)
- Temporal processing campaigns (e.g. iFloodS, IPHEX2014)

What is different than a commodity could?

- Comes close to matching HPC levels of performance
- Node-to-node communication critical – high speed, low latency
- Shared, high performance file system
- Management and rapid provisioning of resources

Science Cloud Architecture

- Agile, high level of support
  - Help desk, system administration, security, applications, etc.
- Storage is 90% full prior to use
- The system owns most of the data
- The users bring their analysis to the data
- Extensible storage; build and expand as needed
- Persistent data services built in virtual machines (VMs), containers, or bare metal
- Create purpose built VMs for specific science projects
- Image management

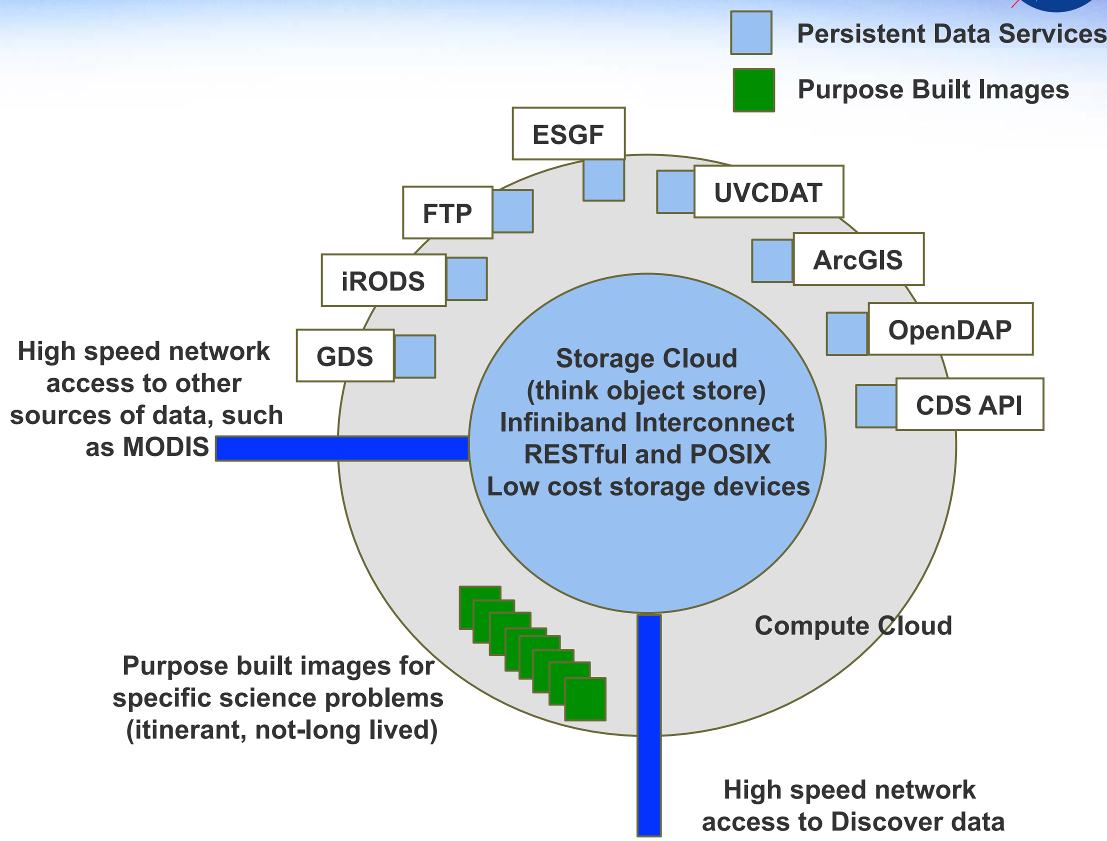

Why not just let scientists use public clouds?

They can!
There are a number of reasons why a NASA managed cloud makes sense…

- Data in Public Clouds
  - Transferring data into the cloud is slow and expensive
  - Storing data in the cloud is expensive
  - Data access in the cloud is not high performance
  - Many copies of data could end up in the cloud multiplying expenses
- Scientists Administering Systems
  - Asking scientists to be their own system administrators
  - Scientists must be knowledgeable in operating systems, tools, license management, file systems, security, etc.
- Low Support
  - There is little to no support in public clouds for science processing
  - Do we really want scientists Google'ing how to install an NFS server in the cloud?

This takes away time and funding from scientists.

With a NASA managed cloud, more science can be done!

Next Steps for the Science Cloud

NCCS Science Cloud

- Prototype in 6 months (August 2014)
- Production in 1 year (March 2015)
- Reuse of NCCS hardware that is being upgraded (normally this hardware would be decommissioned and excessed)
- Purchase of new equipment for redundancy, expanded storage, and additional capabilities

Science Support

- NCCS Data Portal
- ABoVE Mission Support
  Hadoop Cluster as a Service
- Nature Run Data Processing
- Looking for Other Science Opportunities

Special Thanks

THANK YOU

Management and Leadership

- Phil Webster
- John Schnase
- Mark McInerney

The People Who Actually Make Things Work

- Scott Sinno
- Hoot Thompson
- Garrison Vaughan
- Al Settell

The Scientists

- Mark Carroll
- Peter Griffith
- Tatiana Loboda
- Dr. Elizabeth Hoy

And Thank You to the All People That Make This Work!