National Aeronautics and Space Administration

Distributed Computation Resources for Earth System Grid Federation (ESGF)

American Meteorological Society 95th Annual Meeting
Phoenix, AZ
January 2015

Daniel Duffy – daniel.q.duffy@nasa.gov and on Twitter @dqduffy

NASA Center for Climate Simulation (NCCS) – http://www.nccs.nasa.gov

Goddard Space Flight Center (GSFC) – http://www.nasa.gov/centers/goddard/home/

Phoenix (1/7/2015)

I sent to my wife, and she replied "Blah… that is all I have to say to you!"

Earth System Grid Federation (ESGF)

- **International collaboration focusing on serving the World Climate Research Program Coupled Model Intercomparison Project (CMIP)**
- **Data sets include**
  - Intergovernmental Panel on Climate Change (IPCC) AR5 data
  - Observations for Model Intercomparison (OBS4MIPS)
  - Analysis for Model Intercomparison (Ana4MIPS)
- **Gateway to scientific data collections hosted around the world through web portals**
- **Register for access, discover, download, some small amount of data analytics**

The issue is that data holdings are getting to big to just download to scientists workstations or even supercomputing centers!

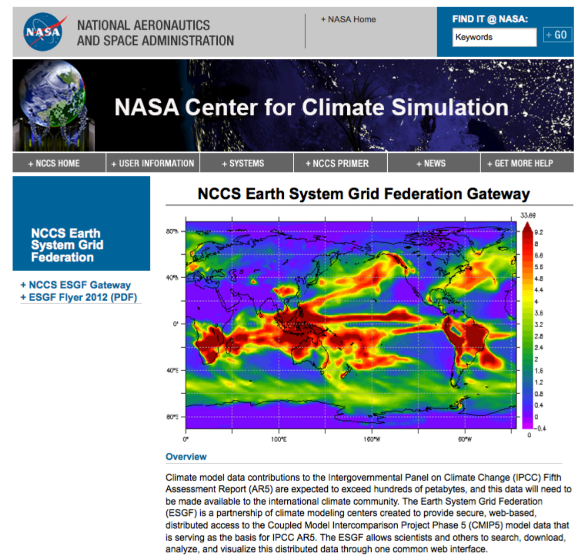

Projecting Future Data Analytics Requirements

How much climate data?

- NASA MERRA Reanalysis Collection ~200 TB
- NASA MERRA2 Reanalysis Collection ~400 TB
- Total data holdings of the NASA Center for Climate Simulation (NCCS) is ~40 PB
- Intergovernmental Panel on Climate Change Fifth Assessment Report (AR5) ~2-5 PB
- AR6 estimated to be 5x(?), 10x(?) AR5
  - Approximate resolution of 25 KM

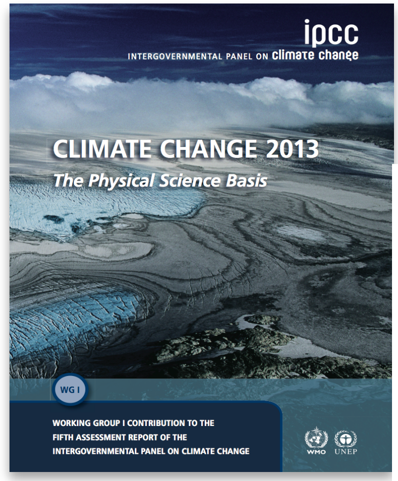

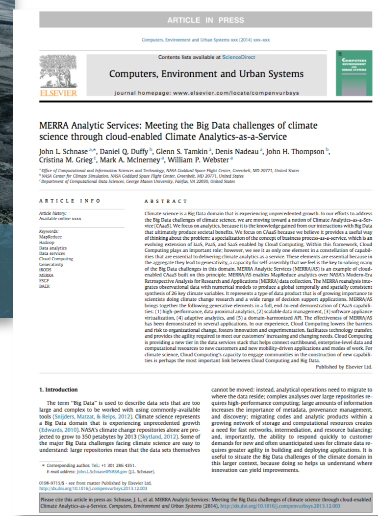
ARTICLE IN PRESS

Computers, Environment and Urban Systems

ELSEVIER

MERRA Analytic Services: Meeting the Big Data challenges of climate science through cloud-enabled Climate Analytics-as-a-Service

John L. Schnase, Daniel Q. Duffy, Glenn S. Tamkin, Denis Nadeau, John H. Thompson, Cristina M. Grieg, Mark A. McInerney, William P. Webster

ARTICLE INFO

ABSTRACT

1. Introduction

Example Problem from Mark Carroll (GSFC)

Change in Surface Water Extent at Beaver Hill Lake Between 2000 - 2010

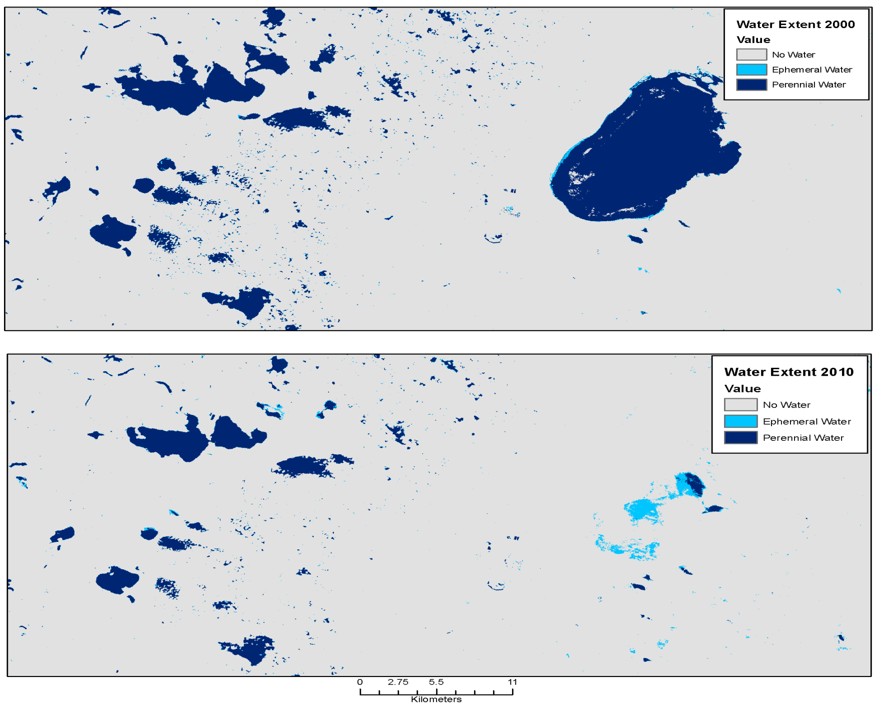

- **Tracking the change in surface water over the study region using Landsat**
  - Average across three epochs (1990, 2000, 2010)
  - 25,000 Landsat scenes/~7 TB of data
  - ***Projected time 9 months!***
- **Download the data into our science cloud**
  - 48 virtual machines
  - 6 weeks of data movement and processing
- **Opened up the opportunity to do more processing**
  - Explore the complete Landsat record
  - 100,000 scenes> 20 TB of data

Climate Analytics as a Service (CAaaS)

Exposure
WPS API

Data

Relevance and Collocation

Data have to be significant, sufficiently complex, and physically or logically co-located to be interesting and useful …

Exposure

Convenient and Extensible and Generative

Capabilities need to be easy to use and facilitate community engagement and adaptive construction …

High-Performance Compute/Storage Fabric

Storage-proximal analytics with simple canonical operations

Data do not move, analyses need horsepower, and leverage requires something akin to an analytical assembly language …

Data

Compute
Storage

Reference:

- Schnase, et al., "MERRA Analytic Services: Meeting the Big Data challenges of climate science through cloud-enabled Climate Analytics-as-a-Service," *Computers, Environment, and Urban Systems* (doi:10.1016/j.compenvurbsys.2013.12.003)

Representative Implementation of a Climate Data Services API at Goddard

QUESTION: Extract the average temperature by season for the year 1980 for the ABoVE region at every level in the MERRA reanalysis data.

```
#!/usr/bin/env python
‘import sys
from CDSLibrary import CDSApi
cds_lib = CDSApi()
service = "MAS"

name       = "above_avg_seasonal_temp_1980_instM_3d_ana_Np"
job        = "&job_name=" + name
collection = "&collection=instM_3d_ana_Np"
request    = "&request=GetVariableBy_TimeRange_SpatialExtent_VerticalExtent"
variable   = "&variable_list=T"
operation  = "&operation=avg"
start      = "&start_date=198001"
end        = "&end_date=198012"
period     = "&avg_period=3"
space      = "&min_lon=-180&min_lat=40&max_lon=-50&max_lat=80"
levels     = "&start_level=1&end_level=42"
file_job_epoch1_aveT  = "./" + name + ".nc"
above_job_epoch1_aveT = job + collection + request + variable + operation + start
+ end + period + space + levels

class UserApp(object):
    if __name__ == '__main__':

        cds_lib.avg(service, above_job_epoch1_aveT, file_job_epoch1_aveT)
```

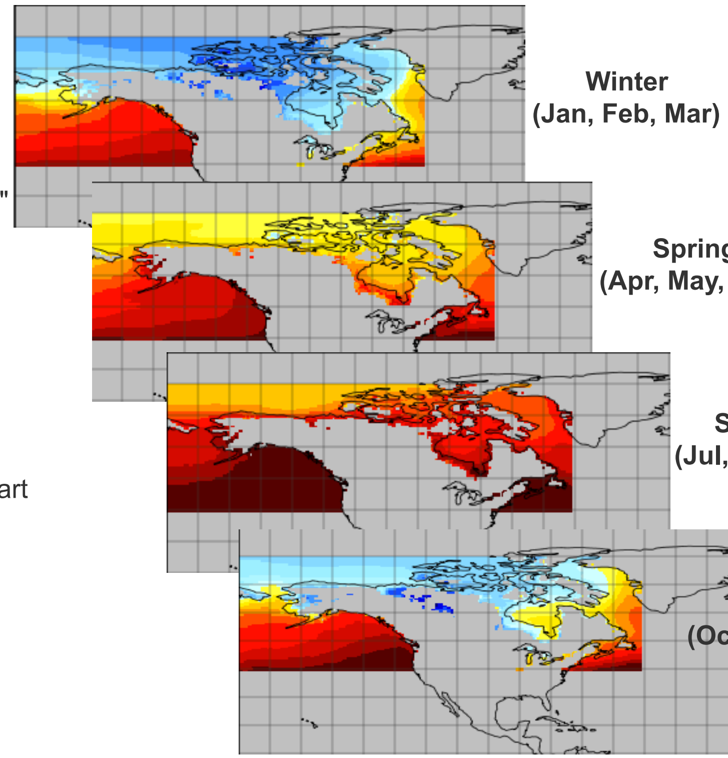

Winter (Jan, Feb, Mar)

Spring (Apr, May, Jun)

Summer (Jul, Aug, Sep)

Fall (Oct, Nov, Dec)

Climate Analytics in ESGF

New methods for climate analytics are needed – access to large geographically distributed data sets and analytics through APIs for storage-proximal processing.

- Creation of the ESGF Compute Working Team (ESGF-CWT).
- Co-chaired by Charles Doutriaux and Daniel Duffy
  - International group of ESGF experts (and me for some reason?)
  - Meet bi-weekly

High Potential for Scientists

- Given that 60% to 80% of a project time is spent on data migration and data wrangling, any speed up will be a great improvement
  - http://earthzine.org/2014/05/22/climate-informatics-human-experts-and-the-end-to-end-system/
- In other words, if it is going to take 9 months for a researcher to get their science completed, we should not worry (at this time) whether or not our server-side capabilities are going to take 1 hour or 1 day
- The bar is pretty low or "laying on the floor" (Eli Dart – ESGF Face-to-Face, 2014)
- Let's get something working and running; we can always optimize later

ESGF-CWT Charge and Progress

Charge of the ESGF-CWT

- Future of ESGF will combine both data AND compute services to create super data nodes
- Develop general APIs for exposing ESGF distributed compute resources (such as clusters, cloud, and/or HPC) to be accessible for large-scale data analytics
- Note that the charge is not to develop the server-side processing capabilities, but rather focus on the API

Process

- Started out with a use case
- Worked through the use case to frame our thoughts and requirements
- Used the Goddard Climate Data Services (CDS) API and server-side processing capability as a driver for the conversation
- Compared and contrasted different viable APIs to settle in on a consensus
    - Settled on the creation of a Web Processing Service (WPS) API

Web Processing Service (WPS)

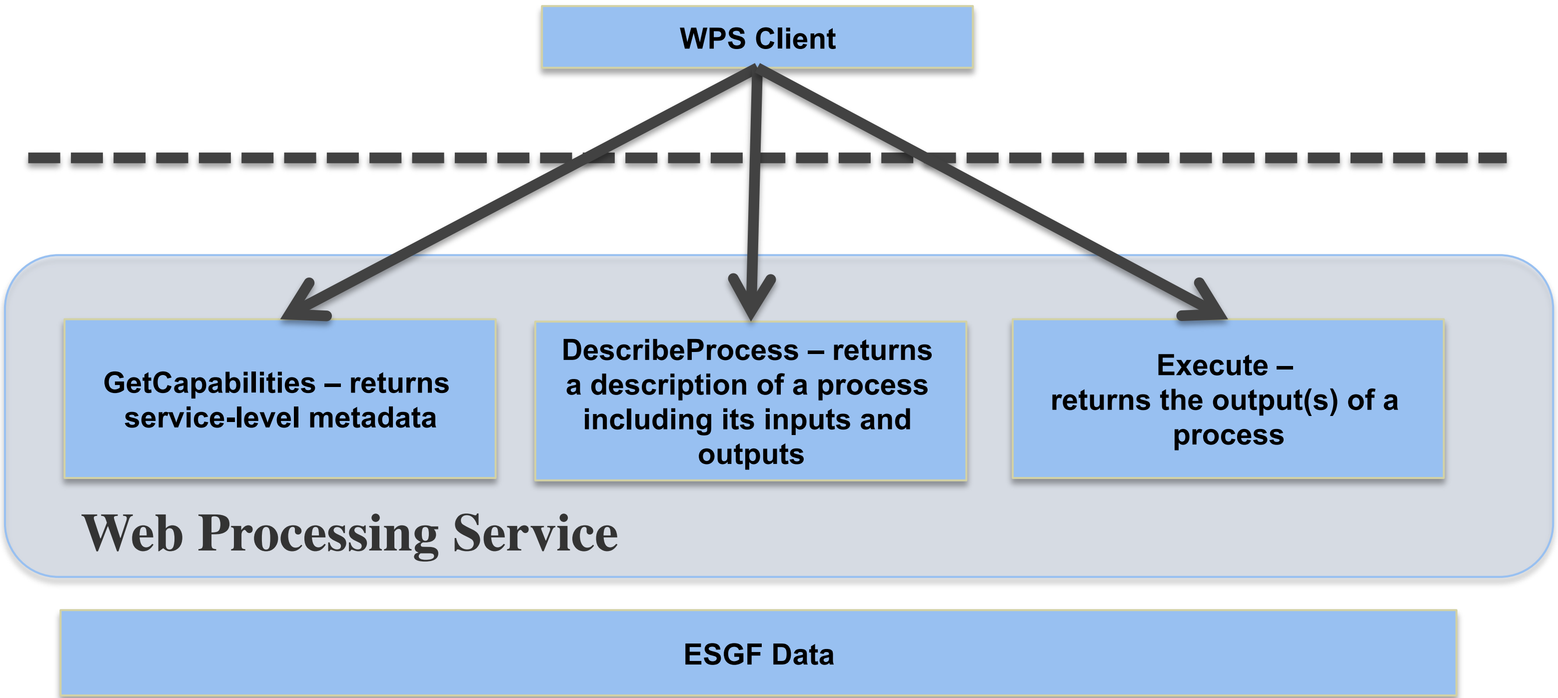

Reference:
- http://en.wikipedia.org/wiki/Web_Processing_Service

ESGF-CWT Representative Use Case – Anomaly

Use Case Statement: Multi-Model Anomaly

- Generate an average of a variable and the difference with respect to that average
- Specify the temporal and spatial extent over which the averaging will be done
- ***Immediate Question:*** How do we handle the grids (space and time) and resolution for each of the models?
  - Use the CMIP5 regridder (ESMF) – resulting steps for the calculation
    - Create the average for each of the individual models.
    - Regrid the models to a common grid.
    - Compute the average across all the regridded data.

Use Reanalysis Data to Work Through This Use Case

- **Representative Questions**
  - Compare and contrast the features from the various reanalysis data sets.
  - Does an ensemble reanalysis provide a more accurate representation of the historical climate?
  - Can we generate uncertainty quantification using multiple reanalysis data sets?

Representative Use Case Workflow Using Reanalysis Data

Work performed by Denis Nadeau (GSFC)

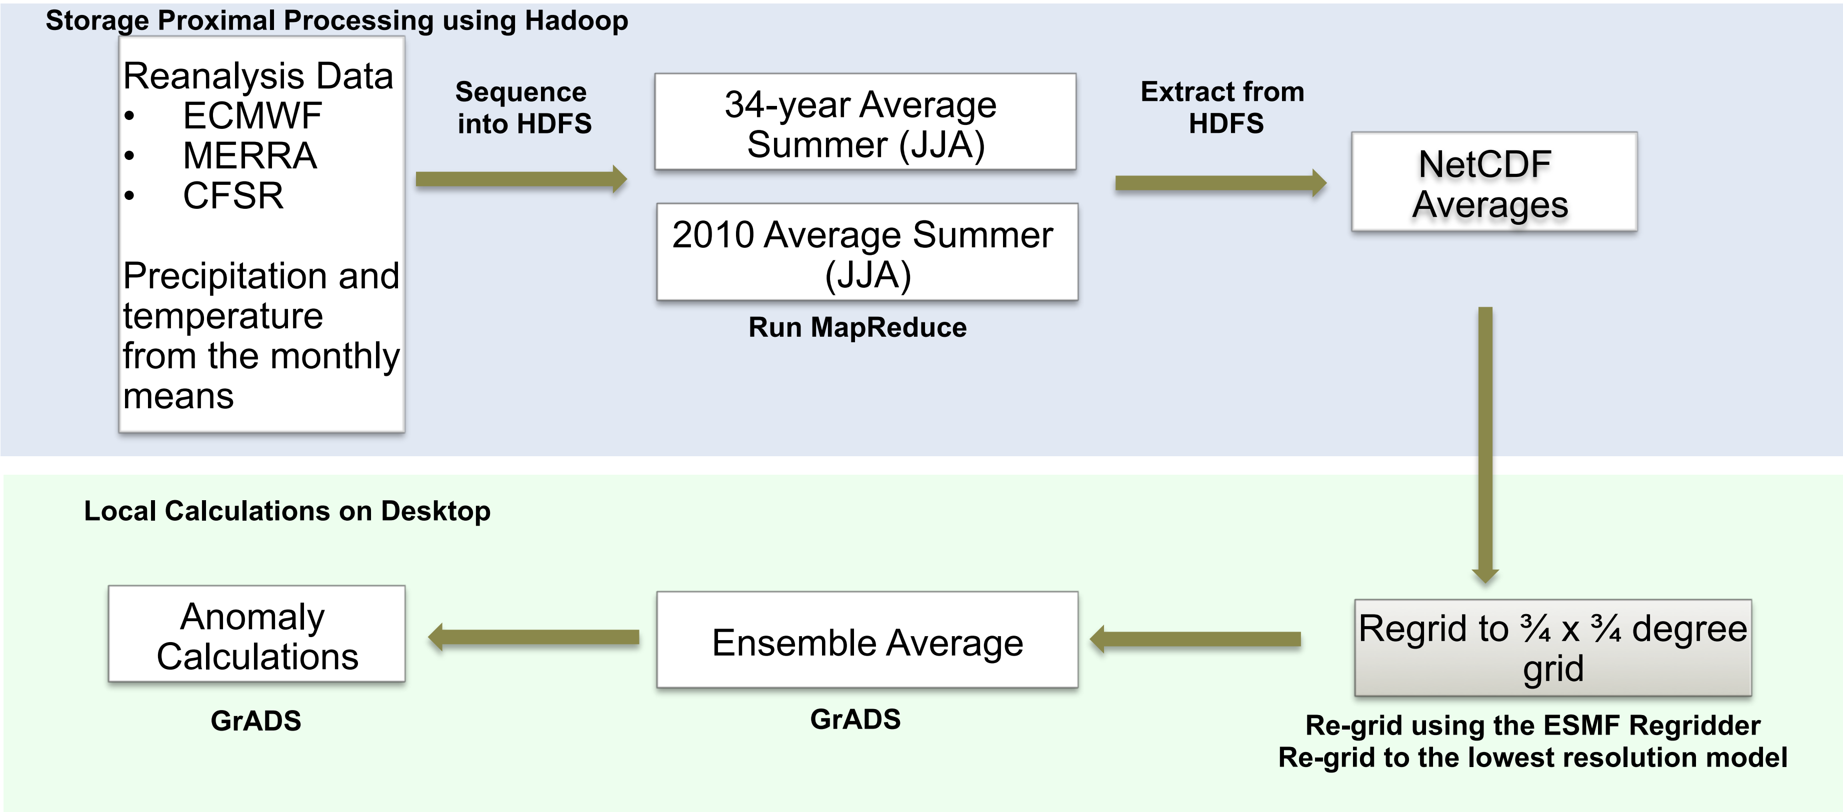

GOAL – Move all of these calculations to the server and expose through an API.

Compute Ensemble Surface Temperature Averages

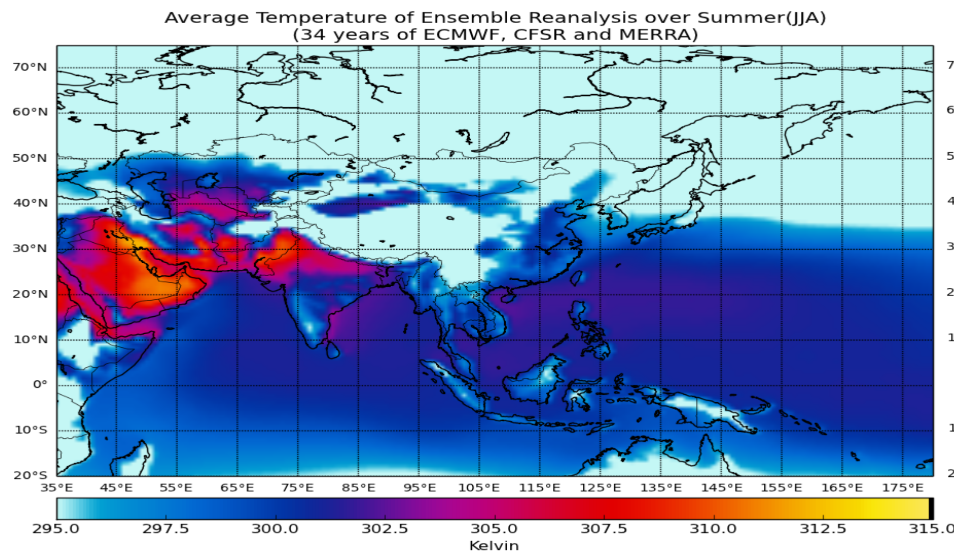

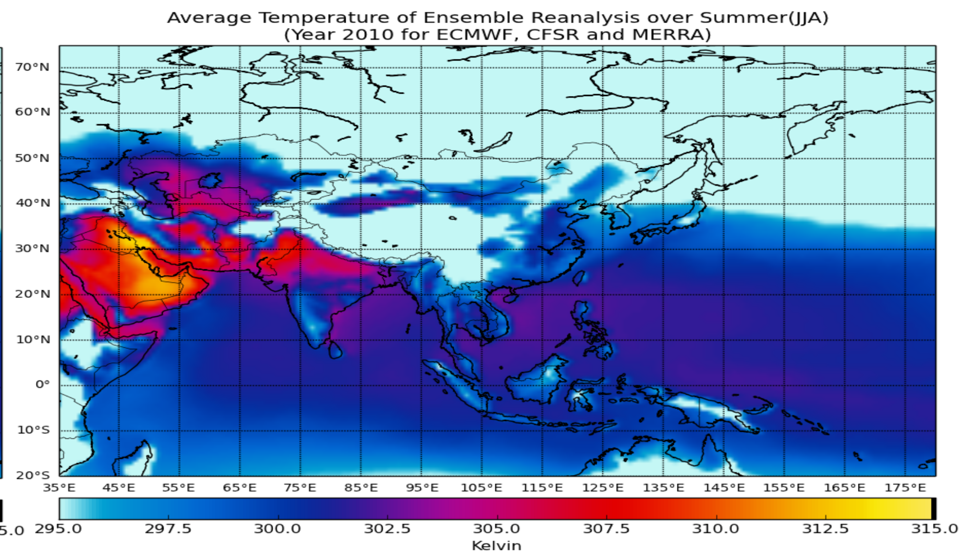

34-Year Ensemble Average over Summer (JJA)

2010 Ensemble Average over Summer (JJA)

Surface Temperature Anomaly of the Ensemble Average

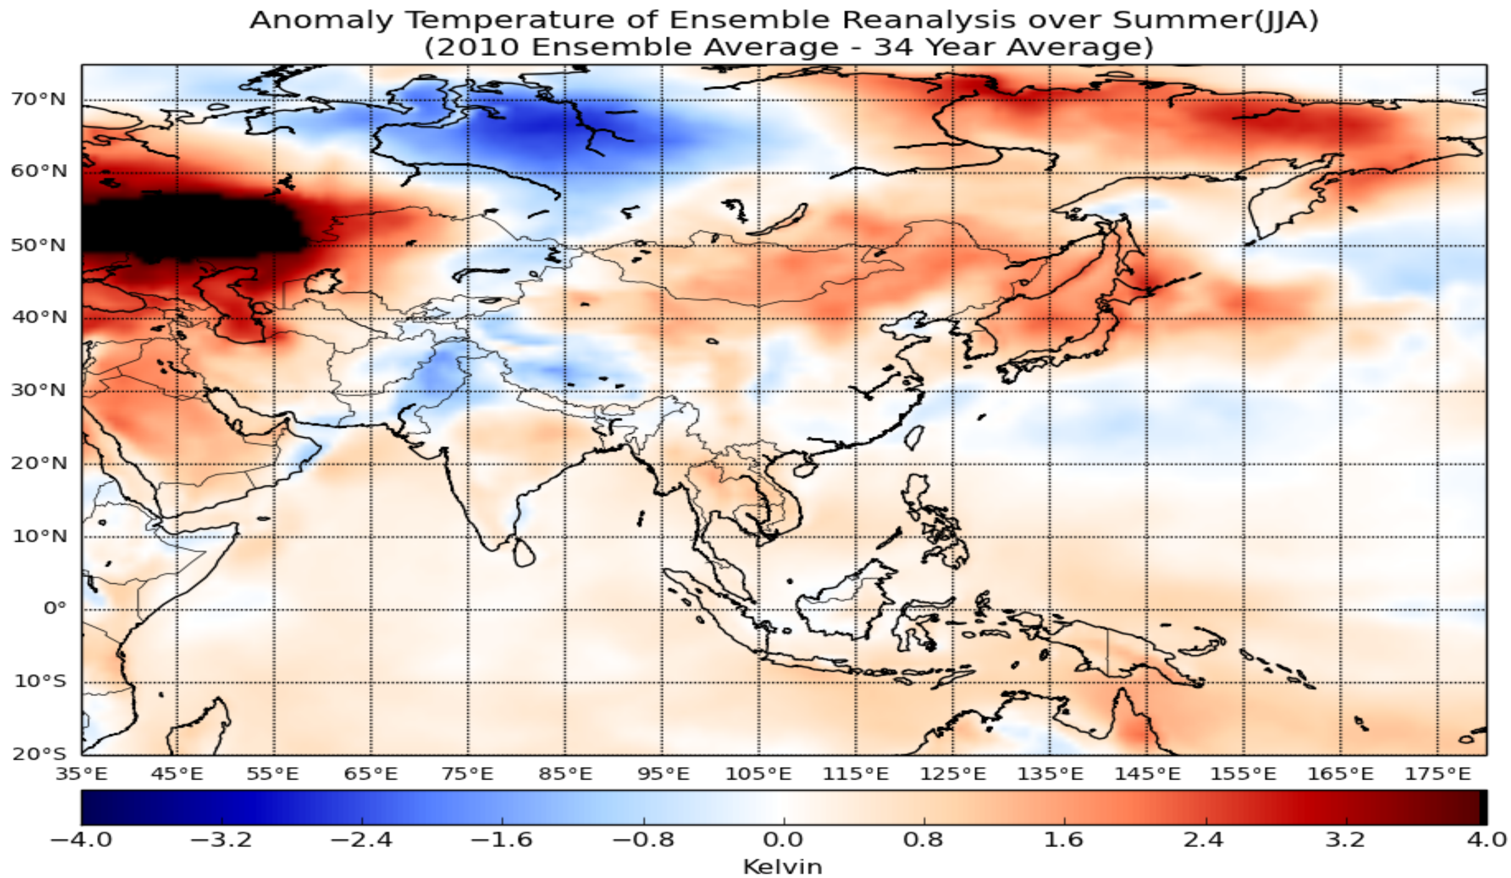

Surface Temperature Anomaly of the Independent Reanalysis Compared to the 34-Year Ensemble Average

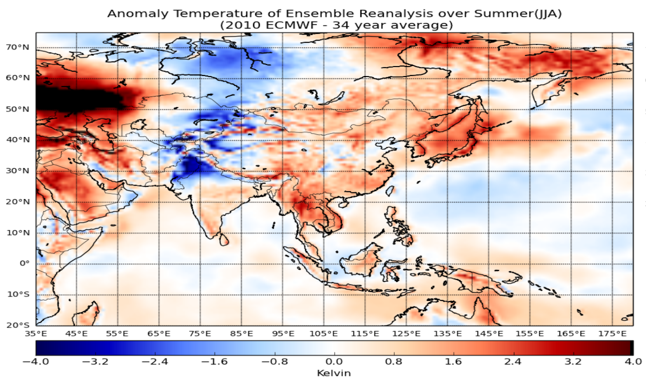

2010 ECMWF – 34 Year Ensemble Average over Summer (JJA)

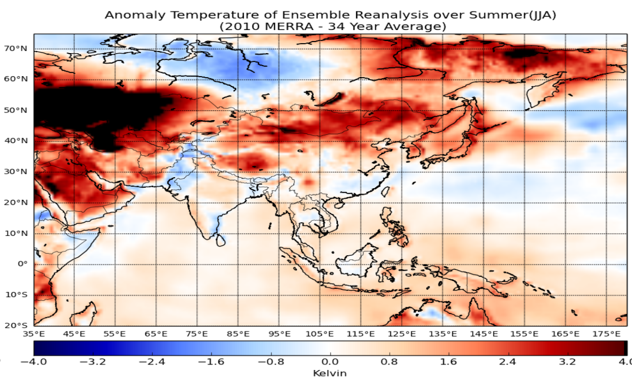

2010 MERRA – 34 Year Ensemble Average over Summer (JJA)

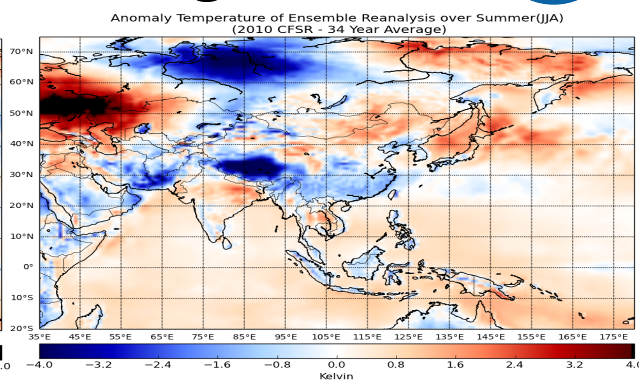

2010 CFSR – 34 Year Ensemble Average over Summer (JJA)

- Temperature variability in the pictures is due to such things as topographical features, different cloud parameterizations within the models, etc.
- Departure of average reanalyses over summer 2010(JJA) shows that ECMWF surface temperature is generally colder than MERRA.
- Himalayas have quite different values!
- Very powerful tool for intercomparisons!

Next Steps Over the Next Year

Data Proximal Analytics Technology at Goddard

- Continue the exploration of HDFS and the ecosystem (specifically Spark)
- Spark (in memory computing)
- Exploration of high performance file systems
- Continue development of the CDS API; compatibility with the ESGF WPS

Application Programming Interface by the ESGF-CWT

- Specification of a WPS API for ESGF
- Reference implementation
- Always looking for more volunteers for this group – contact me if you are interested!

Science

- Compare observations to the ensemble reanalysis
- Work on uncertainty quantifications of reanalysis data
- Engage Goddard scientists to help drive requirements

Future of High Performance Data Storage (Dan's Opinion)

The concept is to both surround and permeate the object storage with compute resources that can be used for analytics.

Tightly coupled HPC systems with relatively large amount of high performance storage

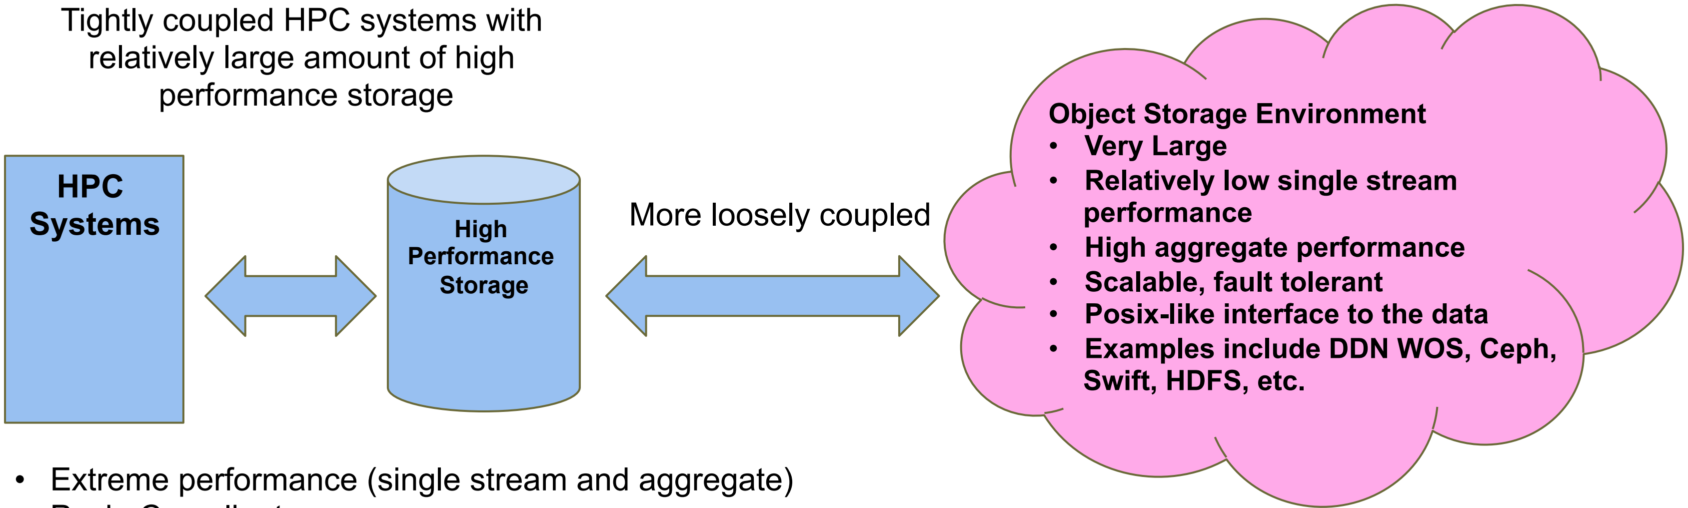

- Extreme performance (single stream and aggregate)
- Posix-Compliant
- Typical examples include GPFS, Lustre, etc.
- Could contain some type of burst buffering capability.

Call it a cloud if you want (or not)! The fact is that we are going to get compute for "free" with the storage in the future, all the way down to the hard drive (check out the Seagate Kinetic Open Storage initiative). We should be working to exploit these capabilities.

Thank You

Very Special Thanks To …

- Funding for this work comes from Dr. Tsengdar Lee and NASA High End Computing Program
- Dean "I Never Sleep" Williams (LLNL)
- Charles "The New Father" Doutriaux (LLNL)
- And all the members of the working group, many of which I met for the first time at the ESGF Face-to-Face meeting in December 2014
- My collaborators at GSFC, including a special thanks to
  - John Schnase, Glenn Tamkin, Denis Nadeau, Hoot Thompson, Garrison Vaughn, Scott Sinno, Dennis Lazar, and all the people at the NCCS