Update on HPC Use for Weather and Climate

Steve Finn Emagine IT

Steve Conway IDC

Introduction

- See Slides from our September meeting

https://hpcuserforum.com/presentations.html

- See videos at

http://insidehpc.com/tag/hpc-user-forum/

- We will focus on 3 of the talks to fit the schedule

Presentations

- High Performance Computing for the Louisiana Coastal Master Plan, Zachary Cobell, ARCADIS-US
  - Use of HPC to model Hurricanes
- Evolution of NASA Earth Science Data Systems in the Era of Big Data, Christopher Lynnes, NASA
  - Observational Satellite data
- NOAA Software Engineering for Novel Architectures (SENA) Project, Leslie Hart NOAA
  - Application software for new architectures

High Performance Computing for the 2017 Louisiana Coastal Master Plan

Z. Cobell

Imagine the result

ARCADIS

Motivation

- Coastal Louisiana is in a land loss crisis.
  - 4800 sq km lost since 1930*
  - 4500 sq km potential loss by 2060*
- More loss, more damage
  - $23 billion additional damage from storm surge due to loss by 2060*

*Peyronnin et al, 2013, Louisiana's 2012 Coastal Master Plan: Overview of a Science-Based and Publicly Informed Decision-Making Process

Agenda

- Numerical Models Employed
- 2007 LACPR
- 2012 Master Plan
- 2017 Master Plan
- Conclusions

Numerical Models

- ADCIRC
  - Circulation model
  - Solves shallow water equations
  - Unstructured finite element mesh
  - Serial or MPI
- SWAN
  - Spectral wave model
  - Solves action density equation for non-stationary waves
  - Unstructured, curvilinear, or structured mesh
  - Serial, OpenMP, or MPI

2007 LACPR

- 2007 FEMA study in Louisiana created suite of 446 synthetic hurricanes
- These 446 storms used to define the statistical returns for water surface elevation

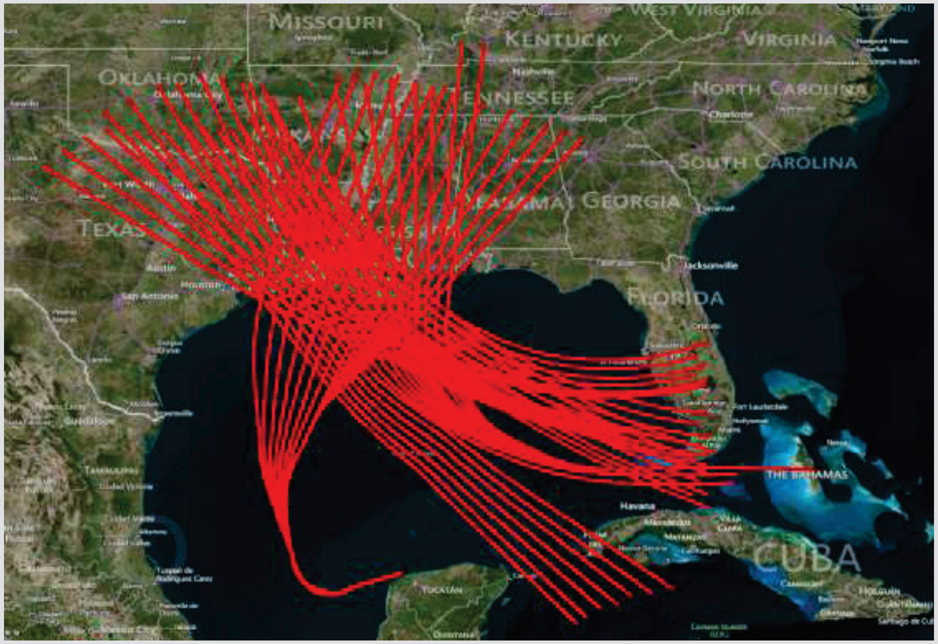

ARCADIS

2007 LACPR

- One of the first studies of its kind
- Developed methods still used today
- Computers simply unable to analyze large numbers of projects
    - Large storm set
    - Computational time required per simulation
- Conducted at USACE supercomputing facilities

2012 Master Plan Framework

- Framework to analyze proposed restoration projects
  - How do these projects effect the landscape?
  - How do these projects effect each other?
  - How do these projects reduce damage during storms?

2012 Master Plan Framework

- Modeling system developed to aid in project selection process
  - Includes models for:
    - Vegetation
    - Water Quality
    - Ecology/Hydrology
    - Wetland Morphology
    - Barrier Island Morphology
    - Fish/Wildlife
    - Storm Surge/Waves
    - Damage/Risk Analysis

Model Interaction

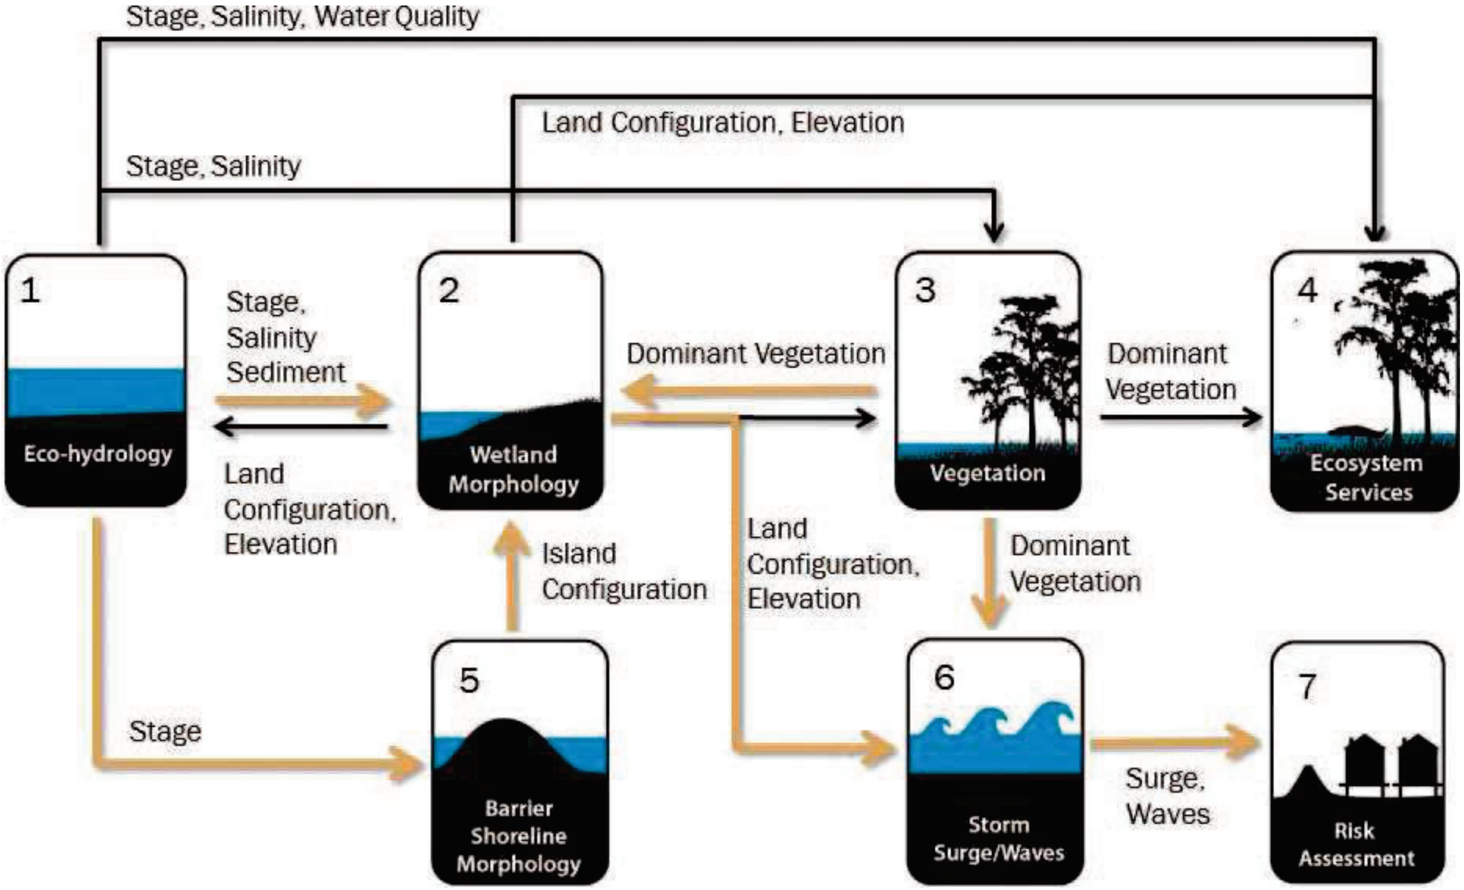

2012 Master Plan

- 2012
  - Total simulations conducted:
    - Project Level

      40 storms
      8 groupings (7+1 base condition)
      3 scenarios

      ---

      840 simulations
    - Alternative Level

      40 storms
      3 groupings
      3 scenarios
      2* time levels (0*, 25, 50 years)

      ---

      800 simulations

1640 simulations

~6MM CPU Hours

2017 Simulation Suite

- 2017 – In progress
  - Number of storms increased
    - 60 storm project screening level suite
      - 3 time levels
      - 3 scenarios
    - 154 storm alternatives analysis suite
      - 3 time levels
      - 3 scenarios
  - Total planned simulations: 7712
  - Total planned CPU hours: ~10MM

Computational Resources

- 2007 LACPR Resource – USACE “Sapphire”
    - Cray XT3
    - Dual AMD Opteron Processors @ 2.6GHz
    - 8192 total cores
    - Shared government resource
- 2012 Master Plan Resource – “Athos”
    - Online 2010
    - Dual 6 core Nehalem Processors @ 2.67GHz
    - Qlogic QDR Infiniband, 2:1 Oversubscription
    - 996 total cores
    - Dedicated resource
- 2017 Master Plan Resource – “Aegaeon”
    - Online 2015
    - Dual 12 core Sandy Bridge Processors @ 2.5GHz
    - Mellanox FDR Infiniband, 2:1 Oversubscription
    - 1488 total cores
    - Dedicated resource

Conclusions

- 2007 – 446 simulations
- 2012 – 1640 simulations (x3.6) – 4MM CPU
- 2017 – 7712 simulations (x4.7) – 8MM CPU
- The number of simulations that can be reasonably completed has grown drastically
  - CPUs, Interconnects, Storage, Availability
- The “value” of a CPU hour has increased
- These simulations heavily influence policy decisions
- The computational power we can harness, the more detail we can provide to policy makers

NASA EOSDIS* Evolution in the BigData Era

*Earth Observing System Data and Information System

Christopher Lynnes
EOSDIS Systems Architect
Code 586, NASA/GSFC
HPC Forum 2015

EOSDIS processes, archives and distributes data from Earth observing satellites

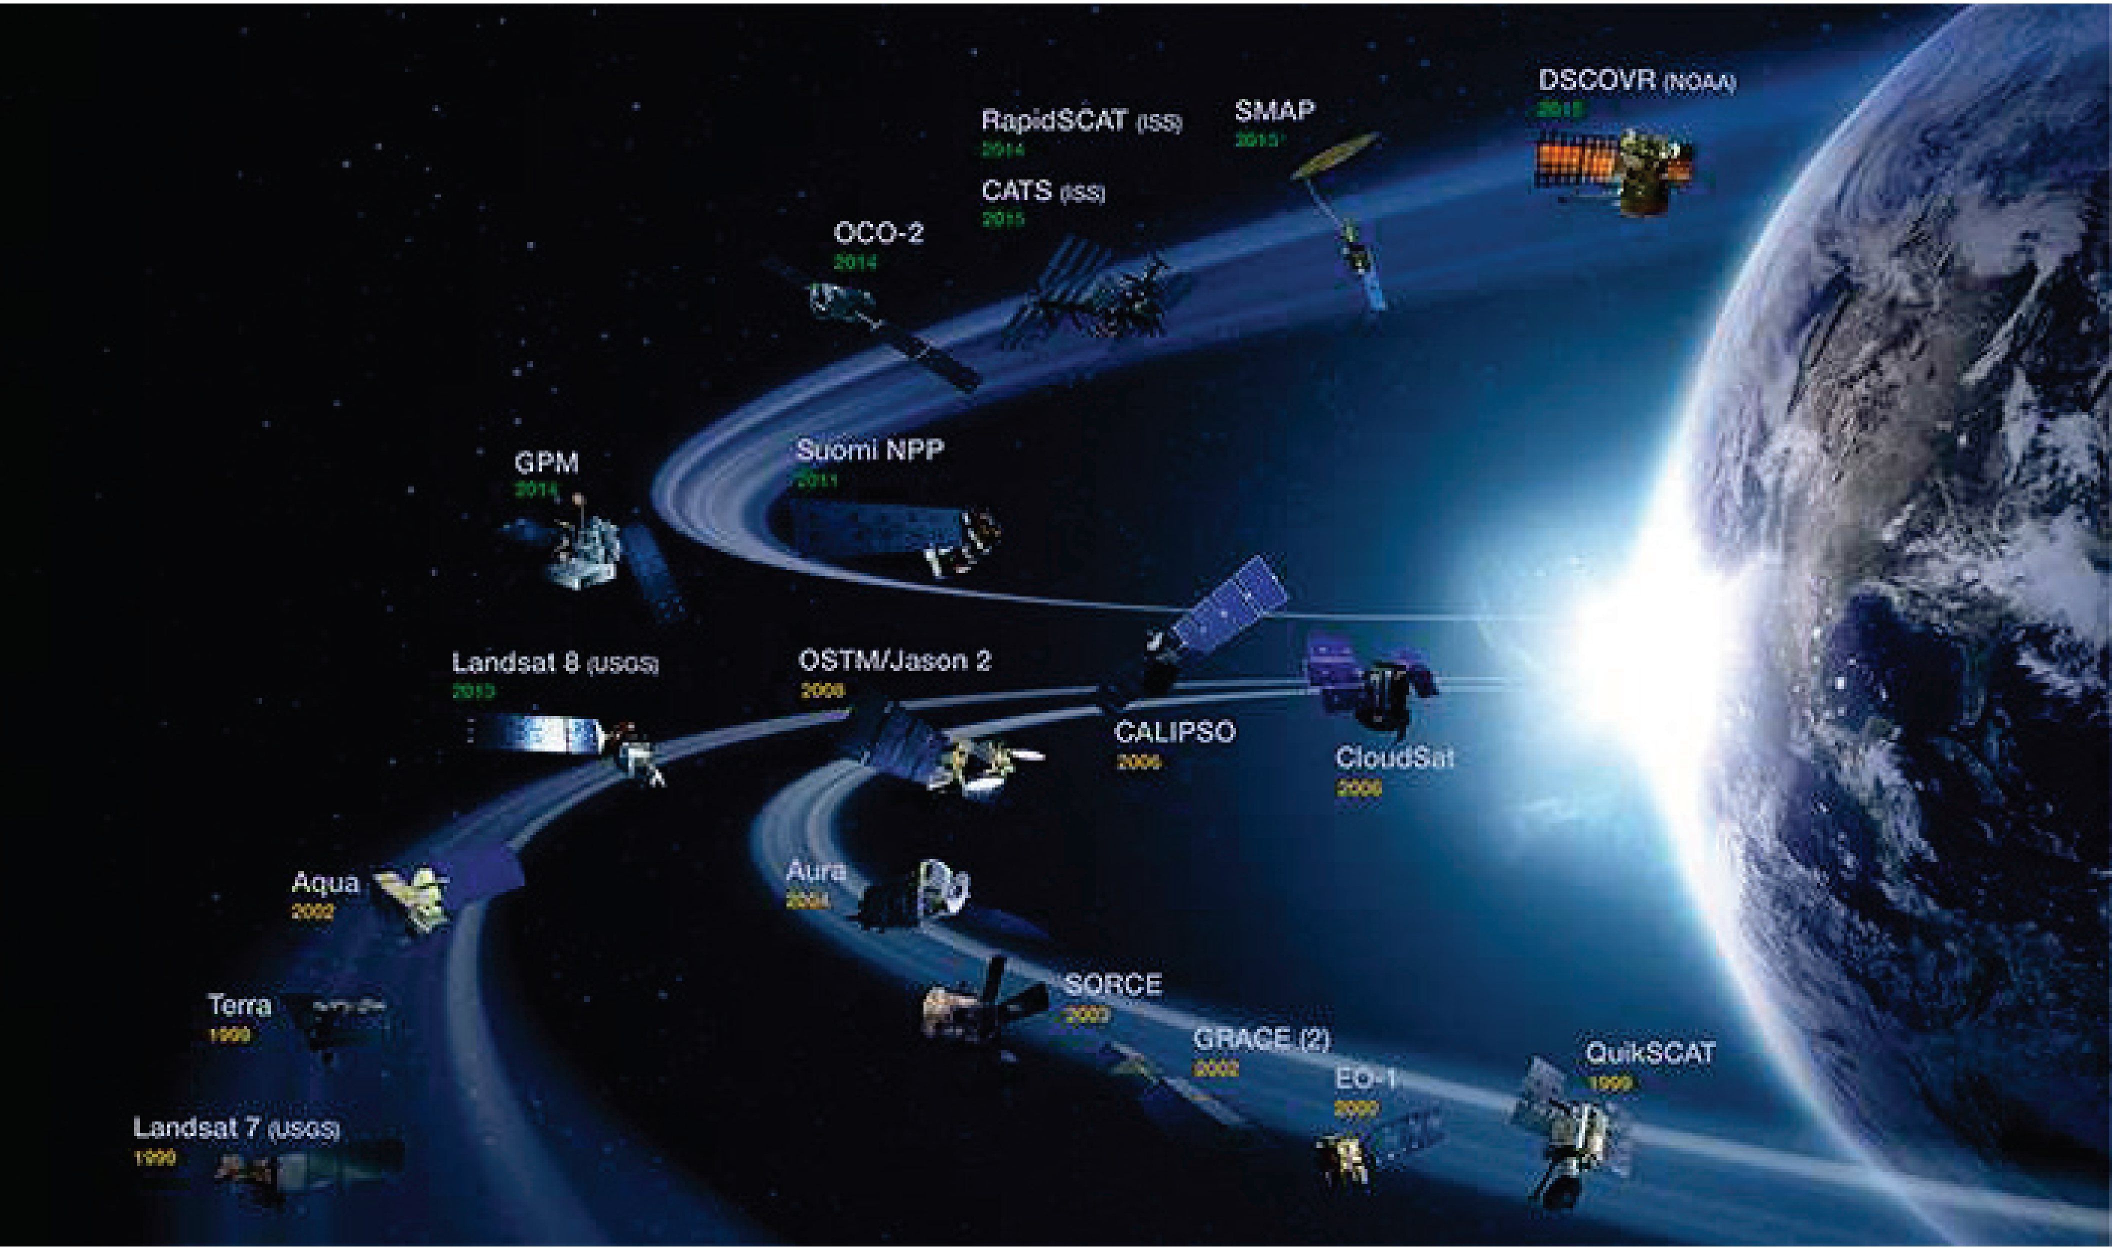

http://earthdata.nasa.gov

EOSDIS manages data from downlink to distribution

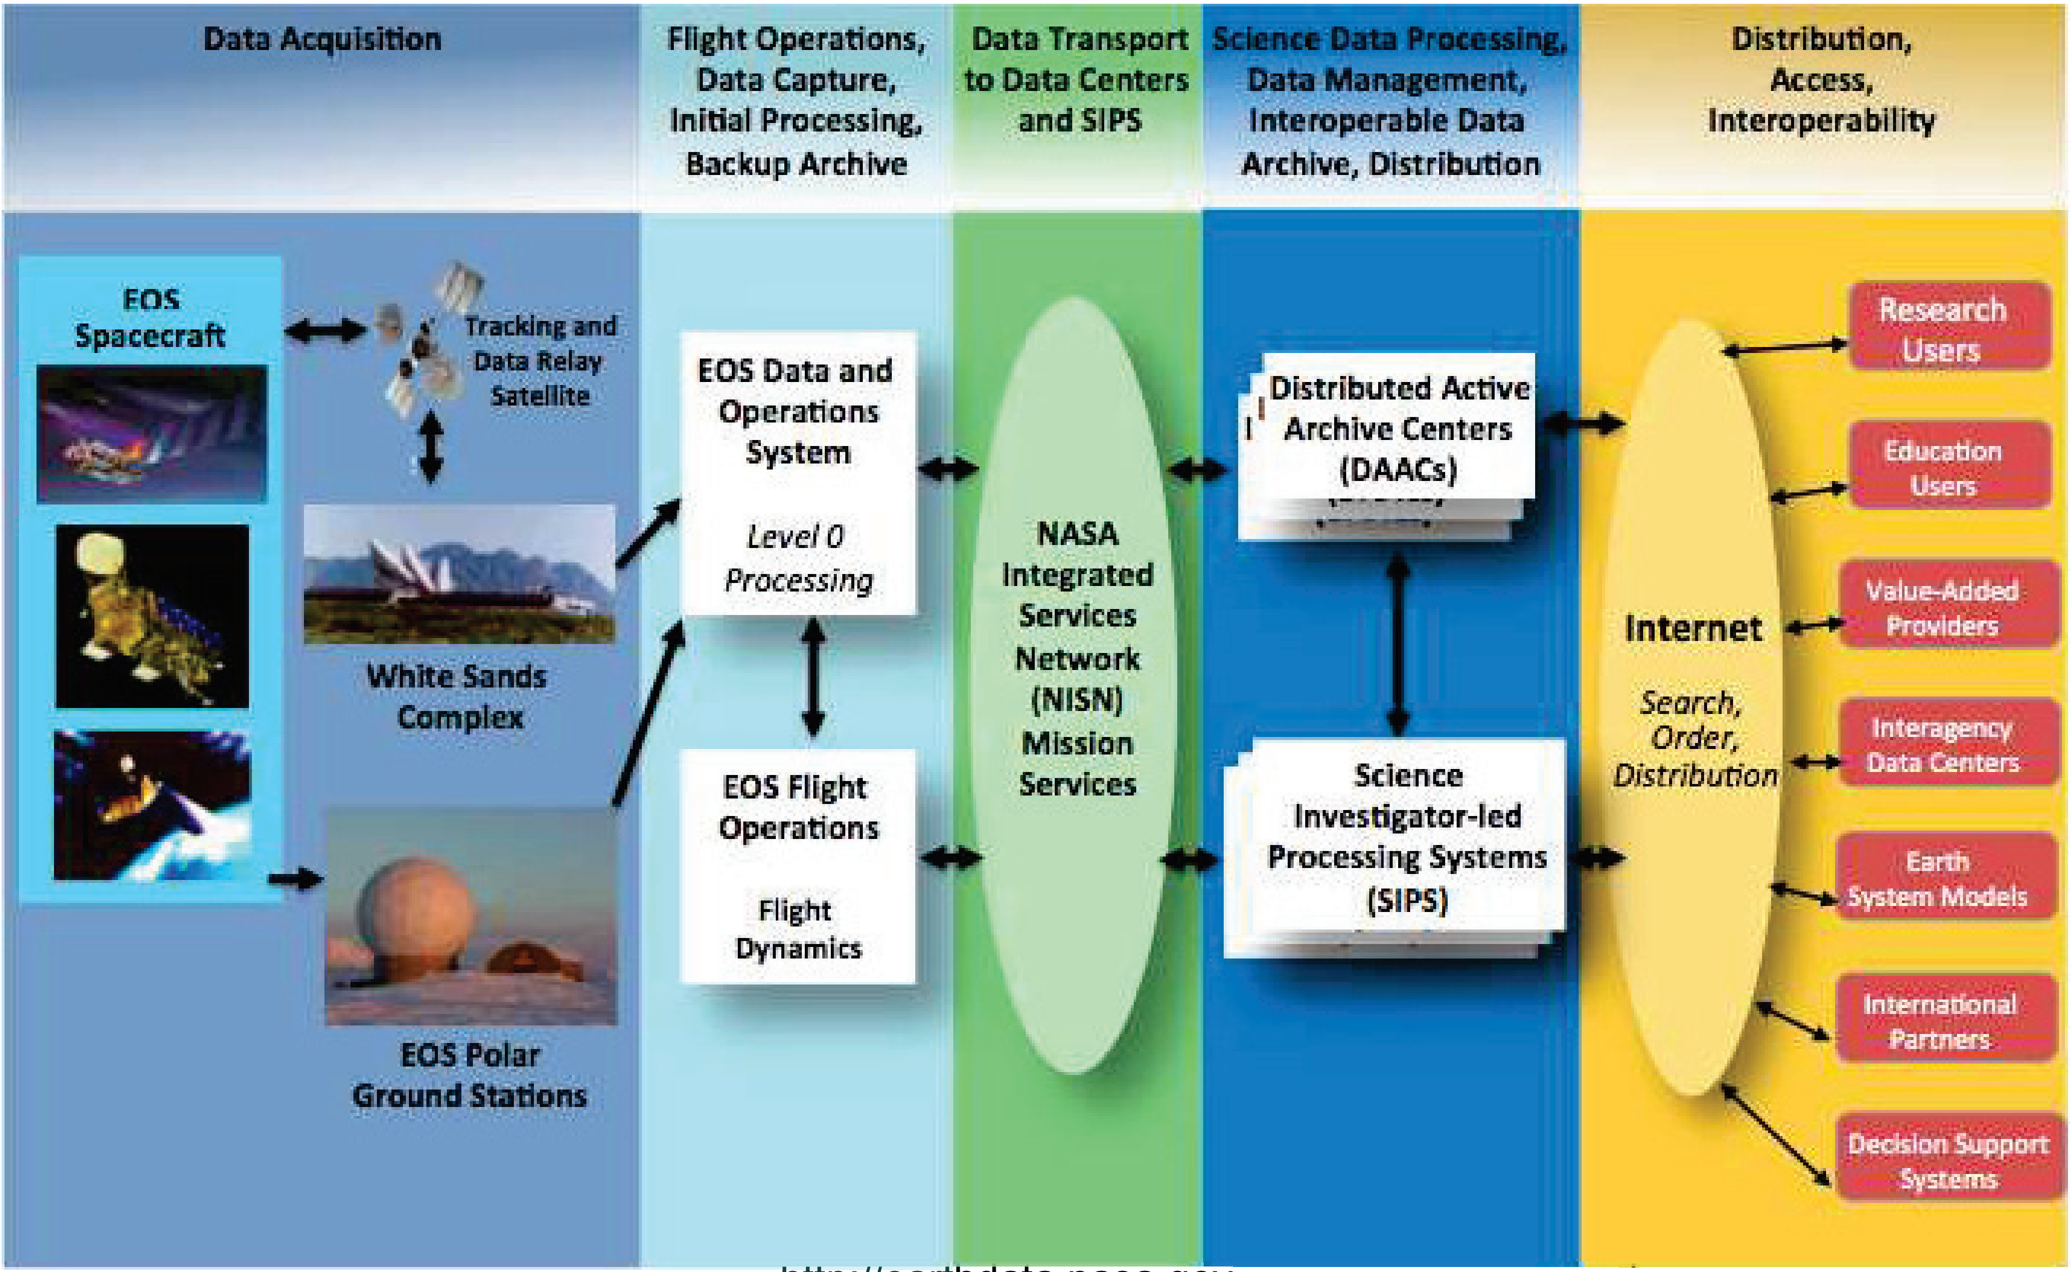

http://earthdata.nasa.gov

DAACs and users are supported by EOSDIS Common Services

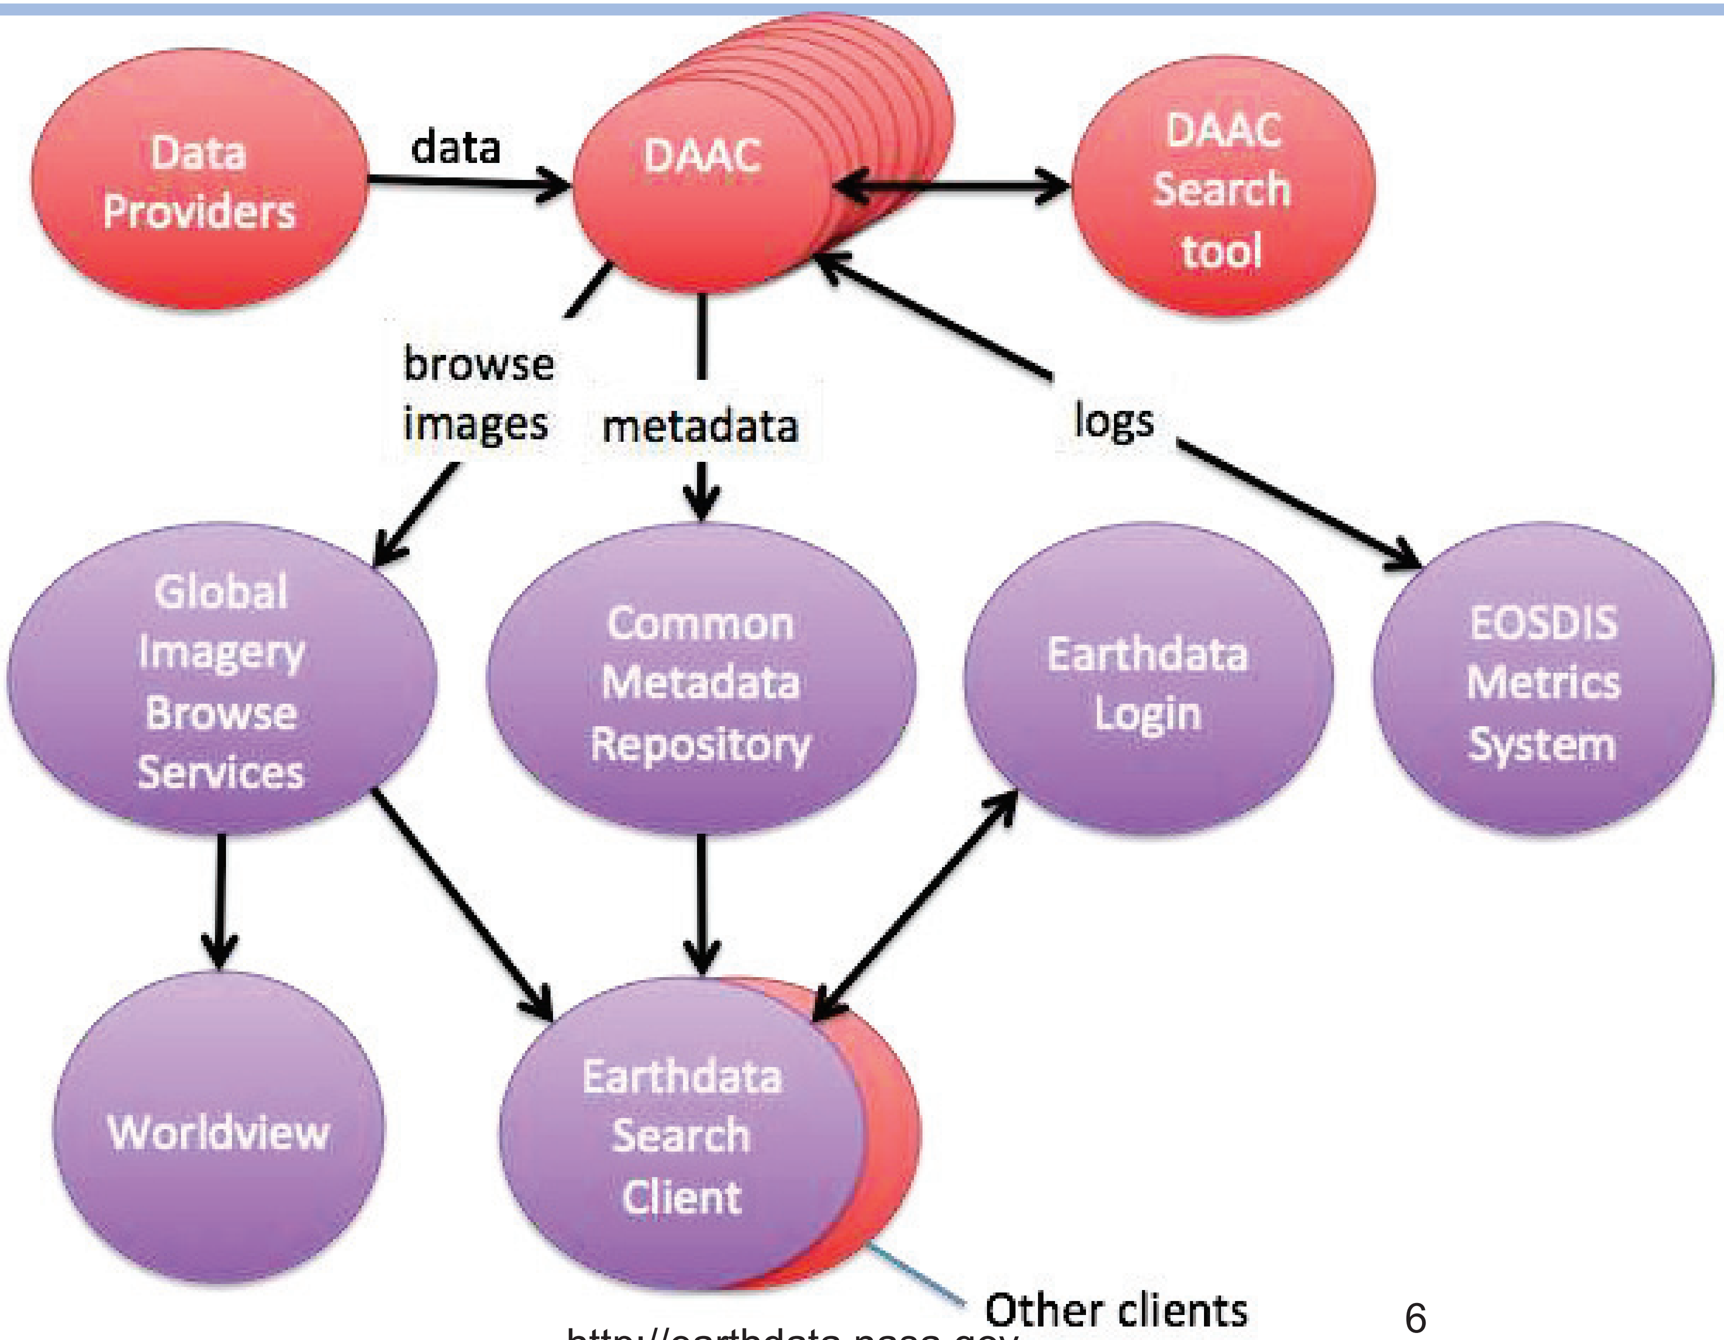

http://earthdata.nasa.gov

EOSDIS Evolves Continually

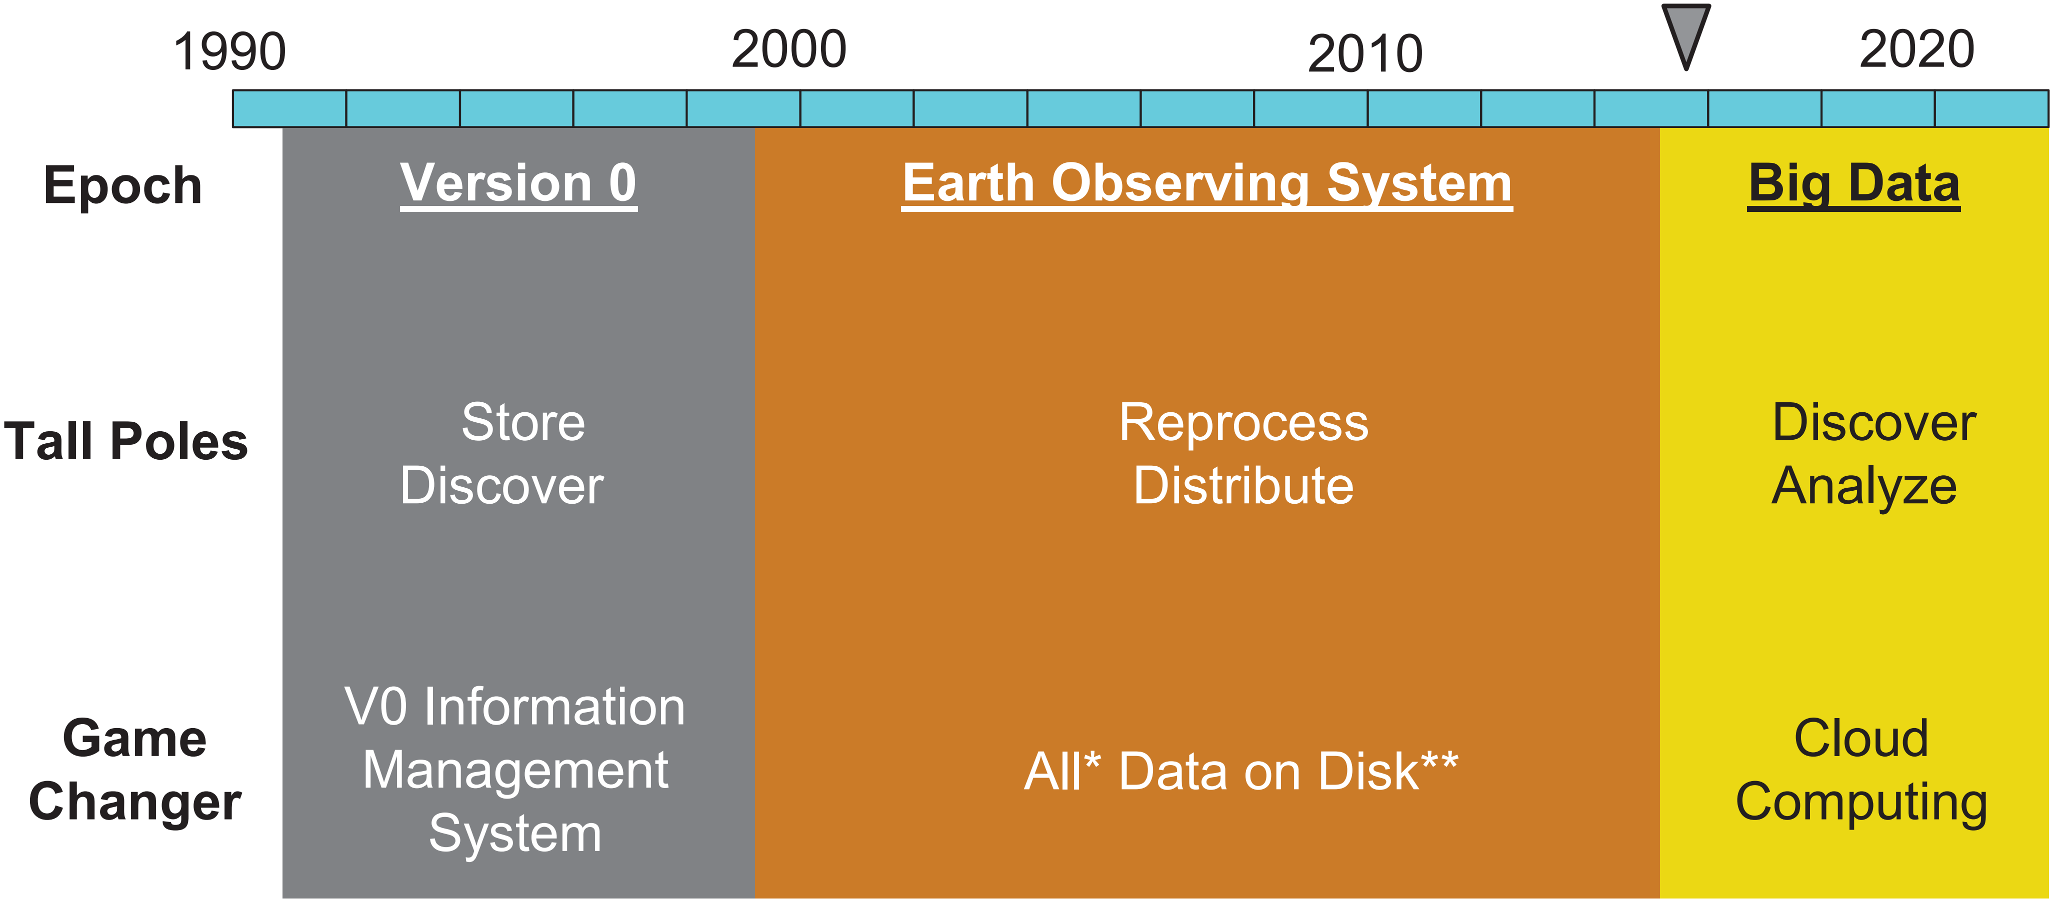

Epoch	**Version 0**	**Earth Observing System**	**Big Data**
Tall Poles	Store Discover	Reprocess Distribute	Discover Analyze
Game Changer	V0 Information Management System	All* Data on Disk**	Cloud Computing

*Almost

**Thank you, HDF internal compression!

Big Data Volume Growth

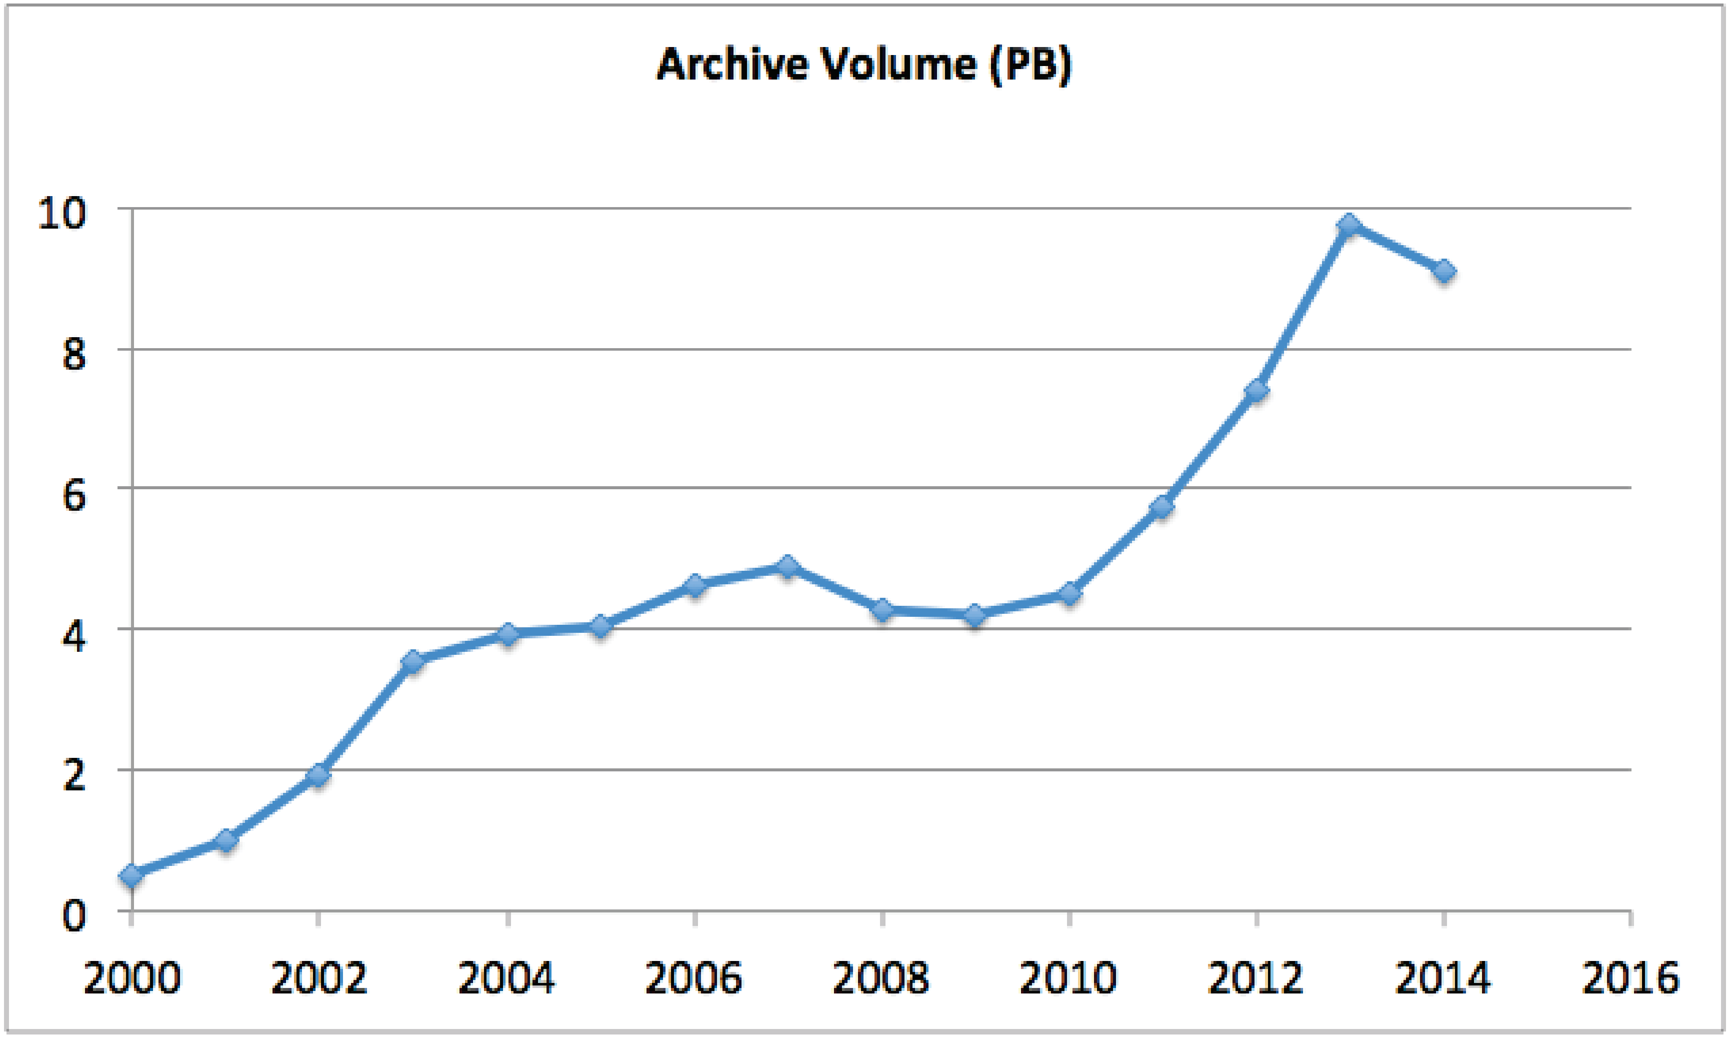

http://earthdata.nasa.gov

Big Data Distribution Growth

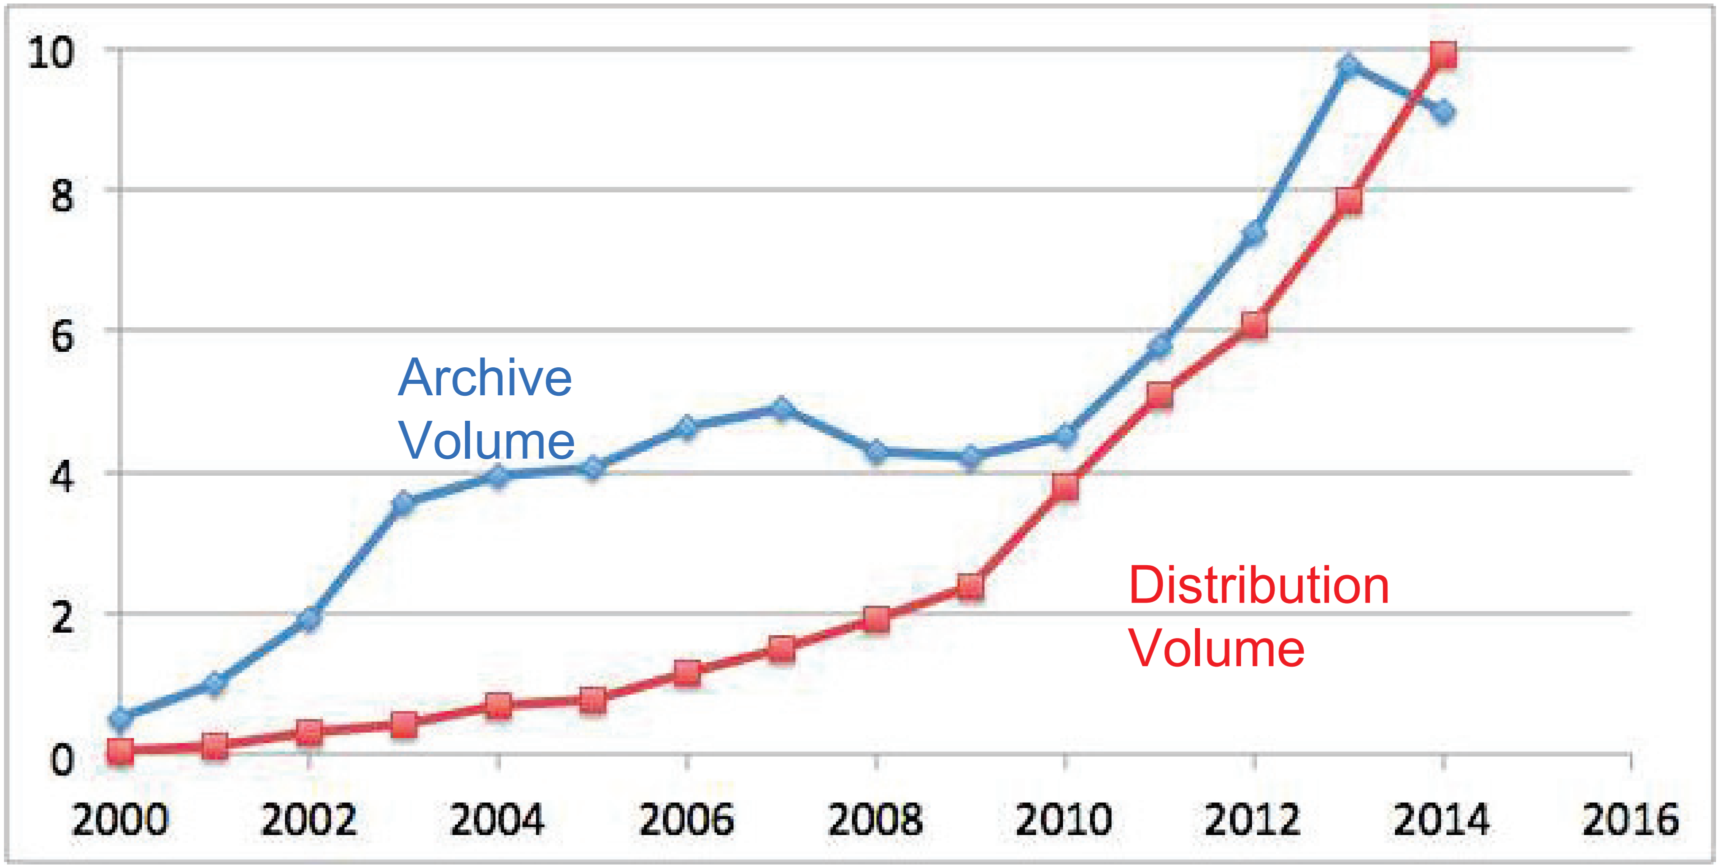

http://earthdata.nasa.gov

EOSDIS in the Big Data epoch will *enable* more analysis closer to the data.

“More Analysis”

More Complexity

Subset

Data Variable

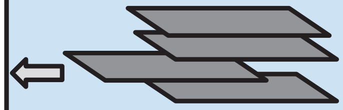

Spatial Area

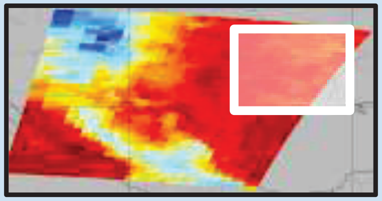

Quality Filter

Transform

Reprojection

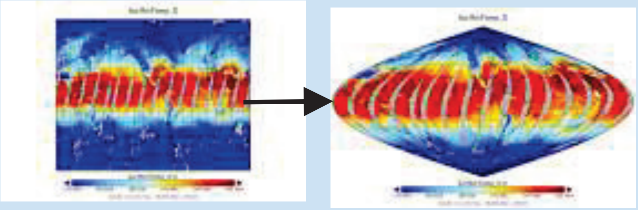

Mosaicking

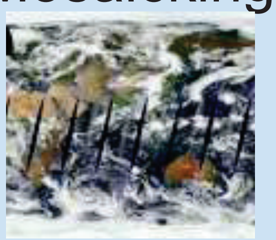

Analyze

Simple Stats

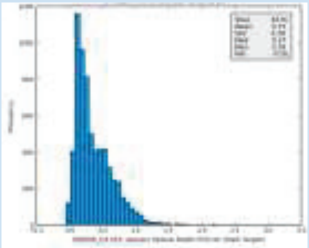

Complex Stats

End User’s Algorithm

```
Welcome to
      ____              __
     / __/__  ___ _____/ /__
    _\ \/ _ \/ _ `/ __/  '_/
   /__ / .__/\_,_/_/ /_/\_\   version 1.4.1
      /_/

Using Python version 2.7.18 (default, May 28
SparkContext available as sc, HiveContext av
>>> 
```

“Close To” = At Archive

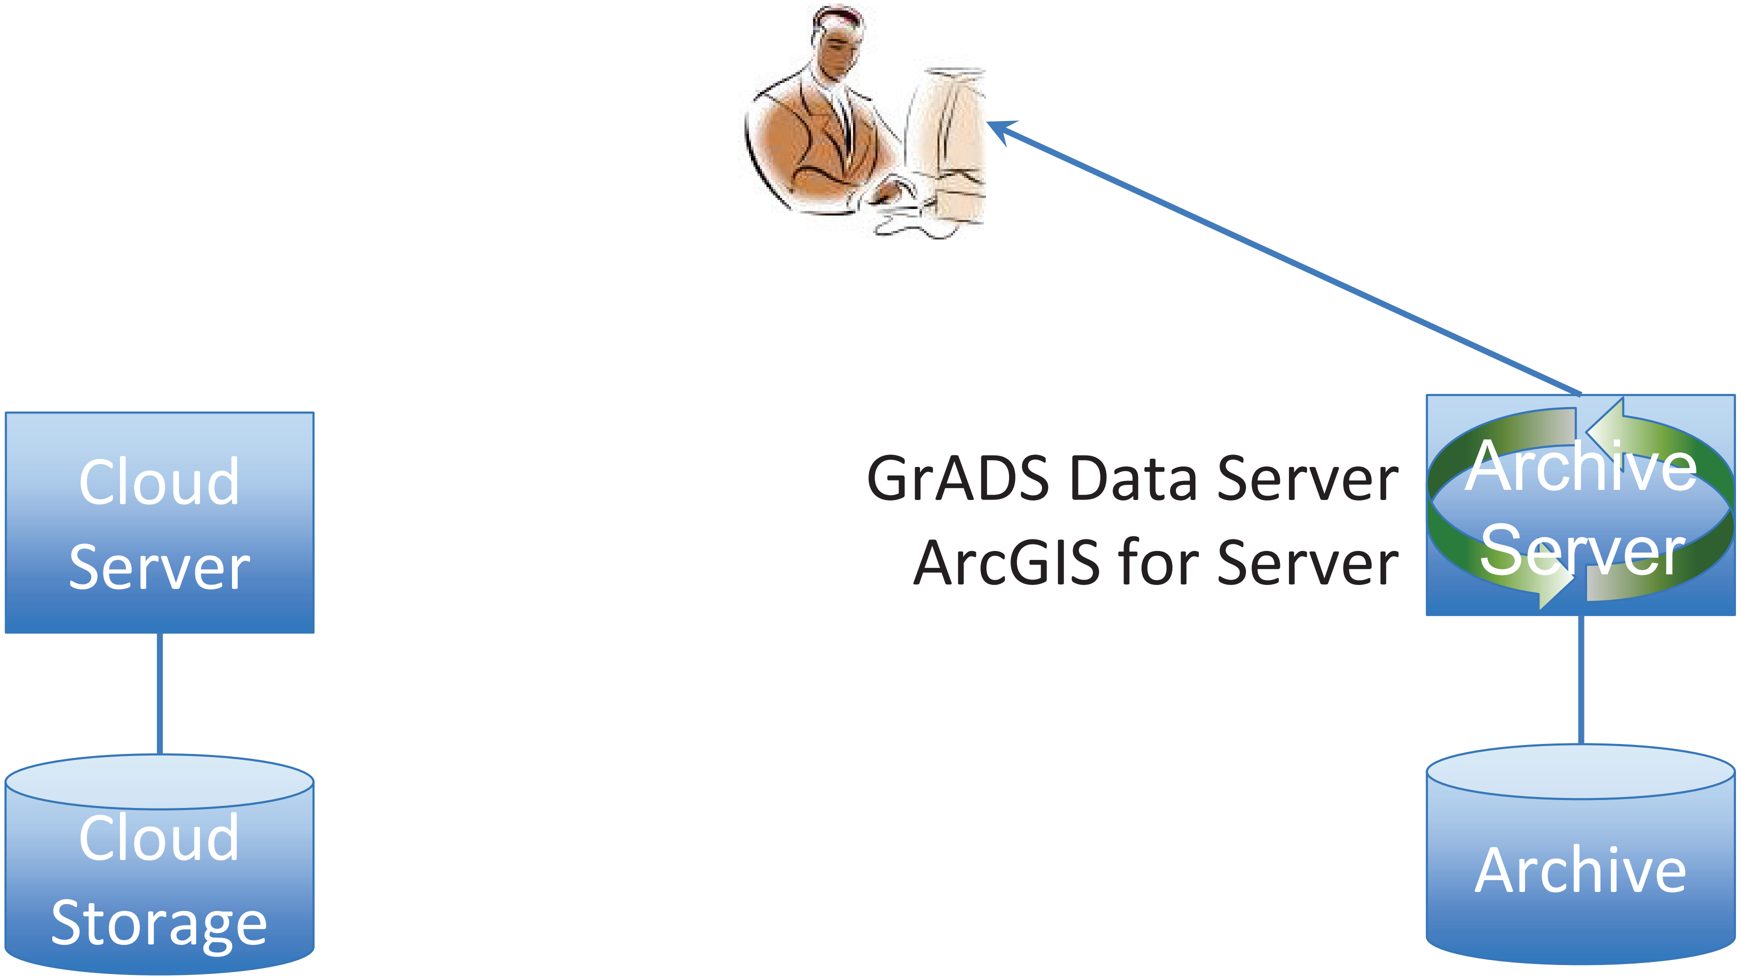

Cloud Server

Cloud Storage

GrADS Data Server
ArcGIS for Server

Archive Server

Archive

“Close To” = Near Archive

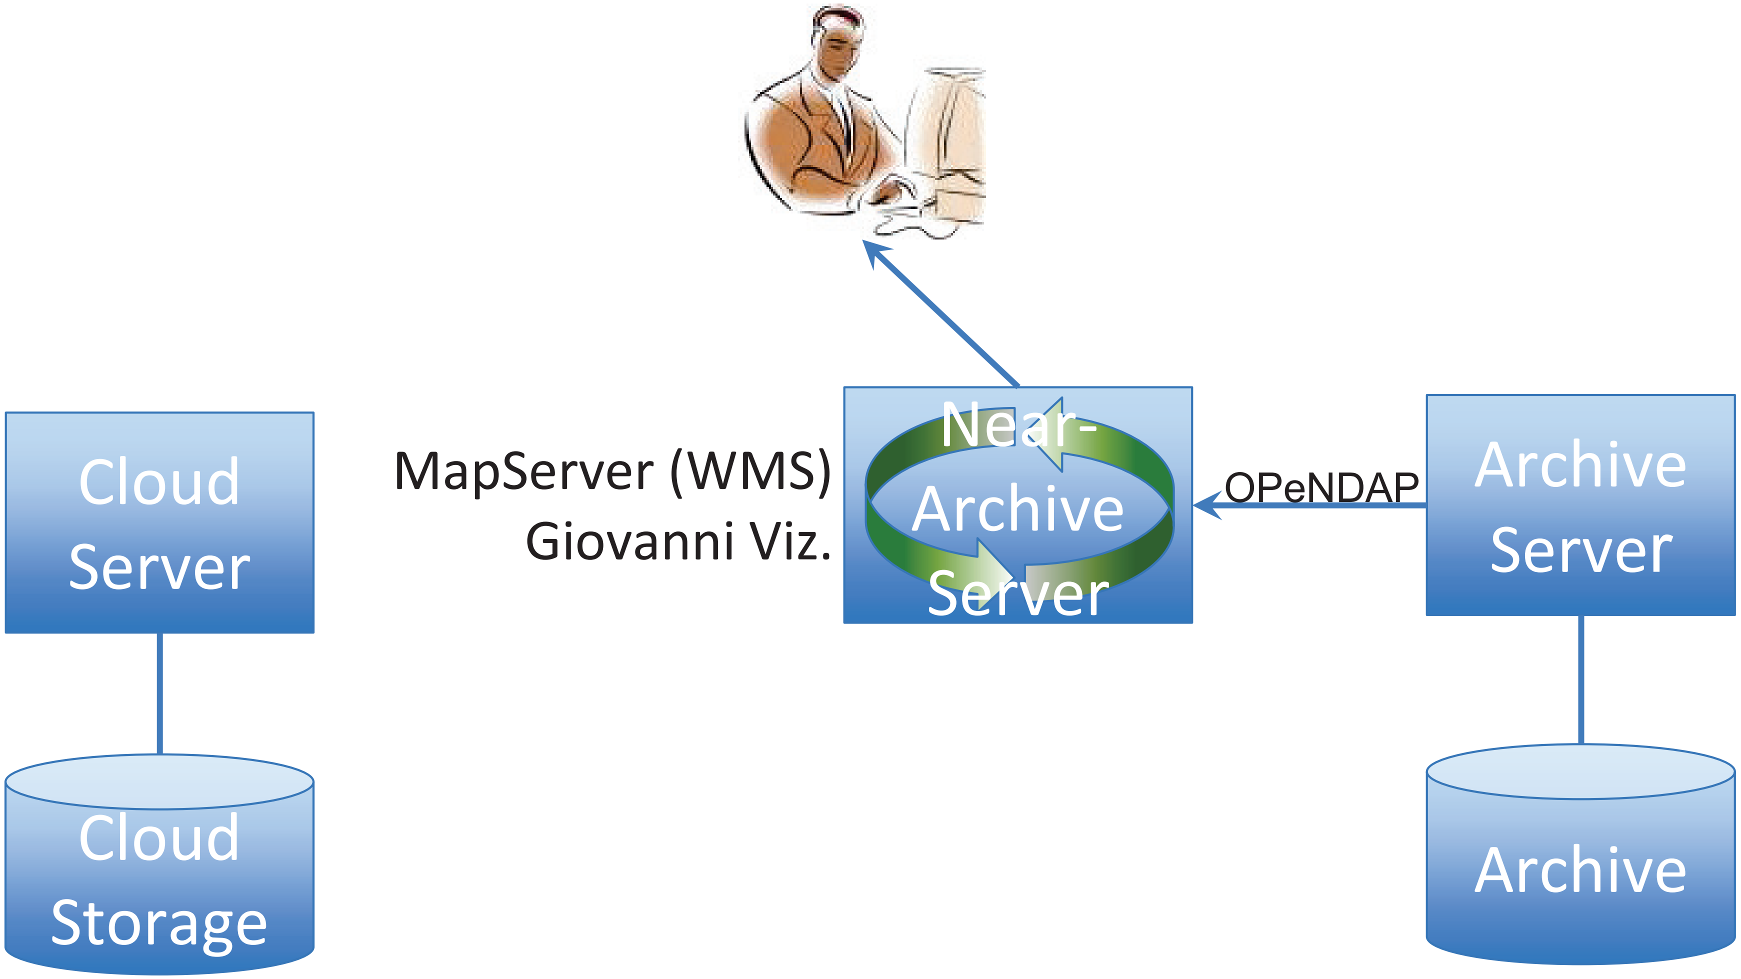

Cloud Server

Cloud Storage

MapServer (WMS)
Giovanni Viz.

Near-Archive Server

OPeNDAP

Archive Server

Archive

“Close To” = Near Processing

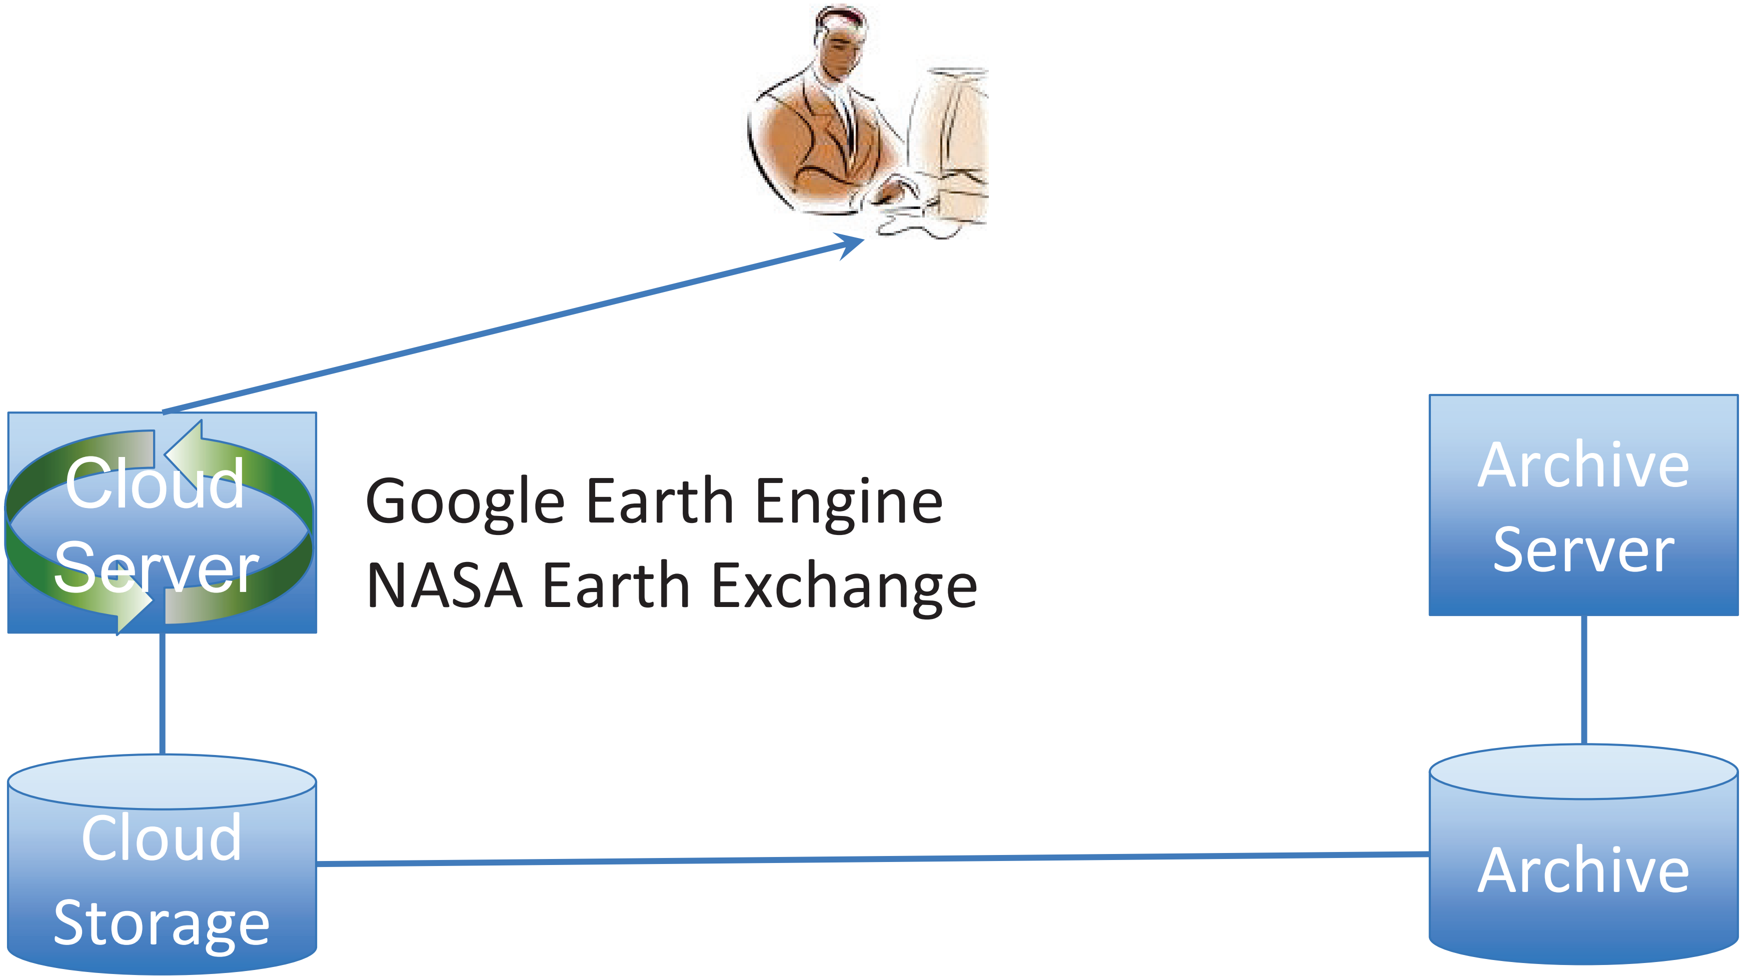

Cloud Server

Google Earth Engine
NASA Earth Exchange

Archive Server

Cloud Storage

Archive

What’s Next?

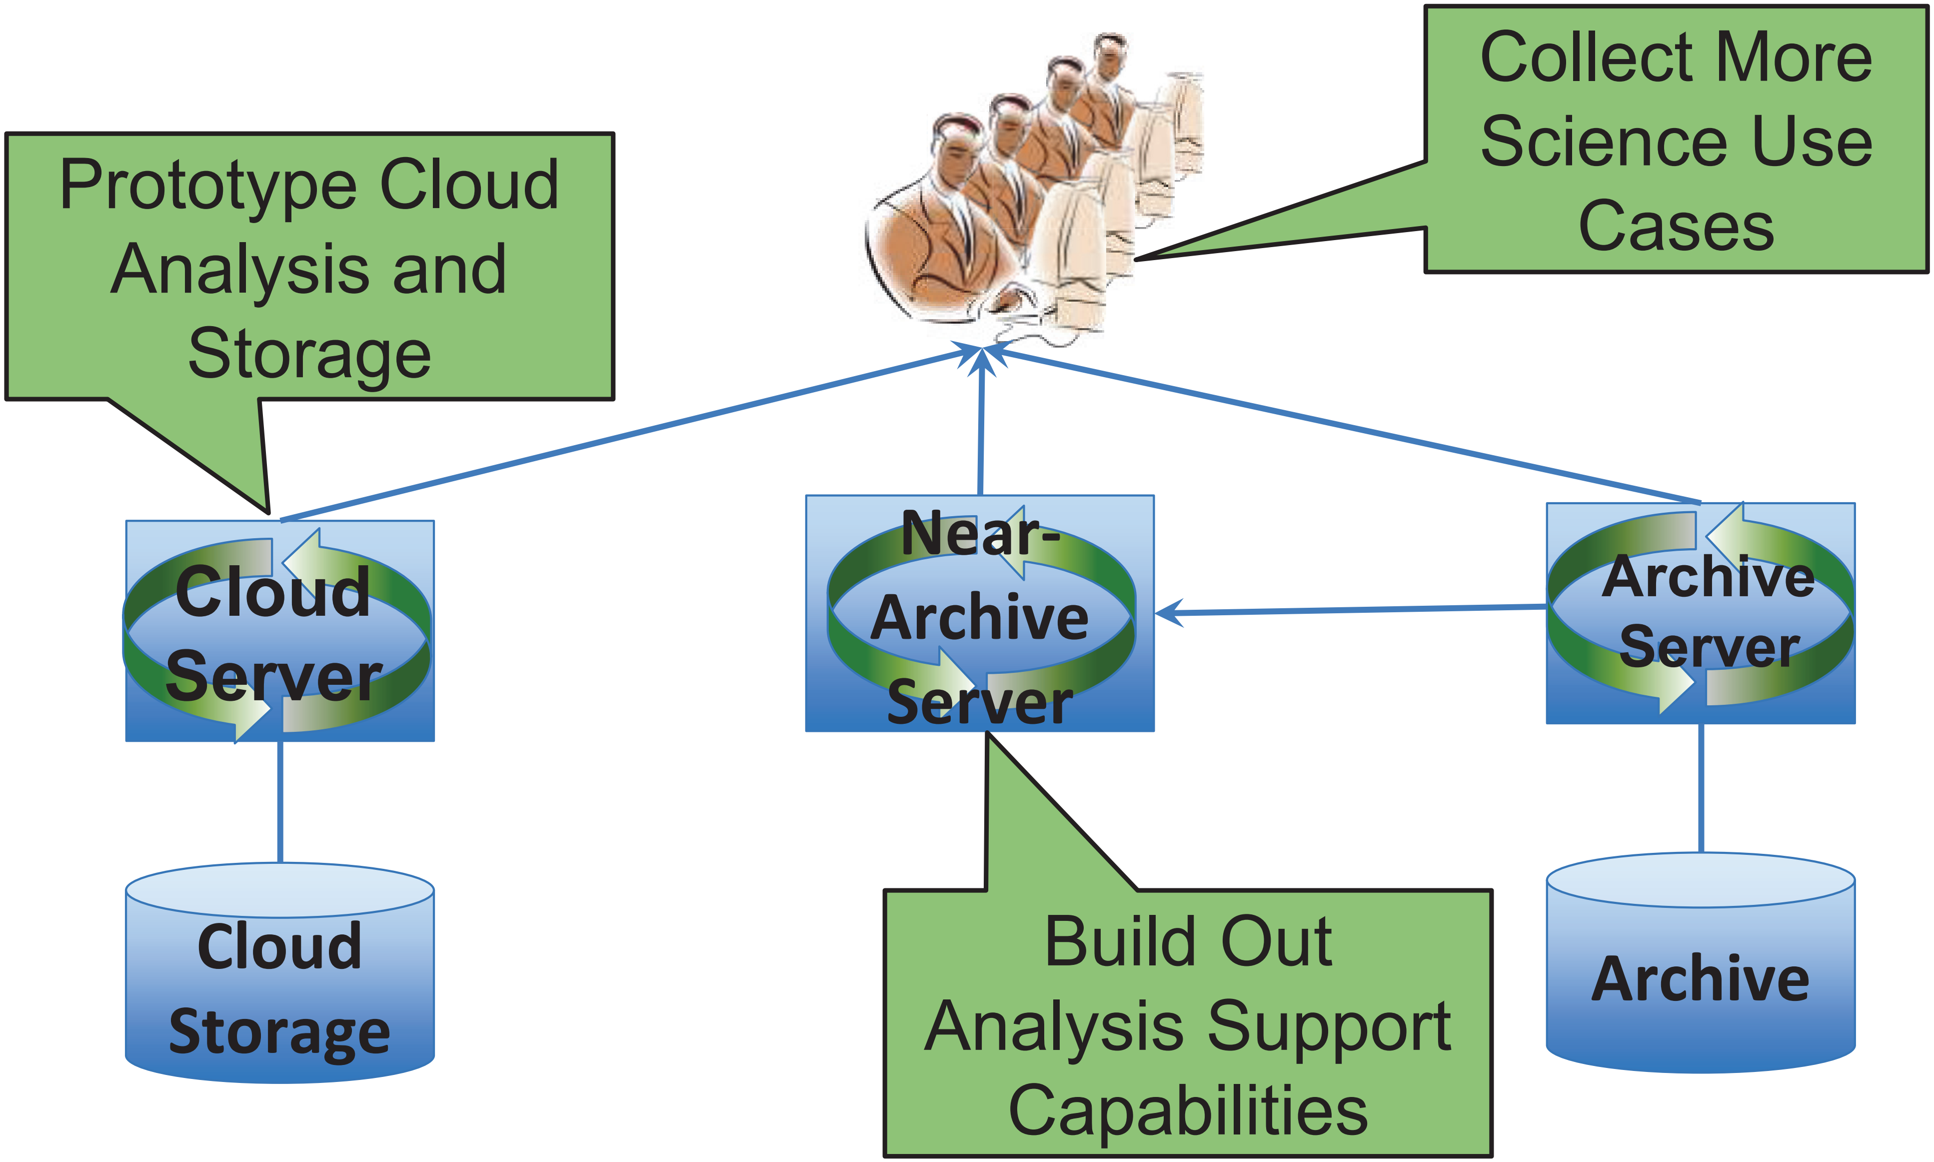

Prototype Cloud Analysis and Storage

Collect More Science Use Cases

Cloud Server

Near-Archive Server

Archive Server

Cloud Storage

Archive

Build Out Analysis Support Capabilities

NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION

Software Engineering for Novel Architectures (SENA)

NOAA Software Engineering for Novel Architectures (SENA) Project

Leslie Hart
9/9/2015

Funding Language

"NOAA will acquire software engineering support and associated tools to re-architect NOAA's applications to run efficiently on next generation ***fine-grain*** HPC architectures."

What is “fine-grain”

From a recent procurement document: “Fine-grain architecture (FGA) is defined as: a processing unit that supports more than 60 concurrent threads in hardware (e.g. GPU or a large core-count device).” *(of course “traditional” architectures are getting to this point as well)*

Overarching Goals

- Prepare codes for a future production architecture
  - Monitor architectural evolution
- Maintain codes in a way that subject matter experts can still modify the code
- Monitor (and participate as appropriate) evolving standards
- Codes should still be viable for current (traditional) architectures
  - It is expected that code optimizations will increase performance on traditional architectures
- Develop expertise within NOAA

Leadership Team - NOAA-wide

- Leslie Hart - OCIO
- V Balaji - GFDL
  - Rusty Benson
- Mark Govett - ESRL/GSD
  - Tom Henderson
- John Michalakes - NCEP

Priorities

Models:

WRF (ARW/NMM), MPAS or FV3, GFDL Climate, NMMB

Programming Research:

Algorithm development, Programming approach

Standards:

OpenACC, OpenMP, LLVM

Challenges

- Lack of standards across divergent architectures
- Large quantity of legacy codes
- Access to developmental platforms
- Uncertainty of performance gains
- Finding qualified staff

Technical Issues

- Language (is Fortran/MPI still the right choice)
- Multiple layers of parallelism
  - Task, Thread, Vector
- Memory footprint of current fine-grain devices
- Atmospheric codes tend to have little data reuse

Implementation

- Develop/cultivate relationships with vendors such as Intel, NVIDIA, etc
- Work with internal tools such as source-to-source translators to implement single source as much as possible
- Work with compiler vendors to ensure directive based approaches are feasible
- Develop small test environment coupled with access to larger government funded machines (DOE, NSF, NASA)
  - Hardware, Compilers, Tools

Status

- Funding distributed to GFDL, ESRL/GSD, NCEP
  - Hiring in process at all three institutions
- Acquiring test systems
- Beginning/continuing work on models
- Evaluating OpenACC and OpenMP versus source-to-source translation

Synergy

- NOAA's Next Generation Global Prediction System (NGGPS) Program
- NSCI
  - NOAA is a "Deployment" agency
    - Develop mission-based HPC requirements to influence the early stages of the design of new HPC systems and will seek viewpoints from the private sector and academia on target HPC requirements *(of course this is a software and hardware problem)*

Conclusion

The SENA project is just starting, it is an attempt to consolidate and accelerate existing NOAA projects

The efforts and results will be useful even in the event that fine-grain architectures do not become as price/performance competitive as some have anticipated

Thank You!