National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Advancing Technology for Big Ocean Science through Partnership and Open Source

Thomas Huang
thomas.huang@jpl.nasa.gov

Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive, Pasadena, CA 91109-8099, U.S.A.

[CL # 19-xxxx]

Cloud: the Brain behind Disruptive Innovations

- "Completely evolved business computing and software architectures, changing a rigid set of services into an iterative, scalable set of applications that constantly transform to meet the need of companies and consumers." – Forbes, Jul 9, 2018
- Power to scale to meet business and consumer needs
- Enables collaboration from different locations
- Streamline delivery of the latest software solutions without shipping boxes of software
- Simplifies access to data and services from any platform
- High availability

- **2018 IDG Cloud Computing Study**
  - 77% of enterprises have at least one application or a portion of their enterprise computing infrastructure in the cloud
  - 76% of enterprises are looking to cloud applications and platforms to accelerate IT service delivery
    - Improving the speed of IT service delivery
    - Increasing flexibility to react to changing market conditions
    - Enabling business continuity
    - Improving customer support and services
  - 95% of all organizations will be relying on the SaaS model for application delivery in 18 months, with IaaS increasing to 83% and PaaS, 73%

Processors are not Getting Faster

2004: First Pentium 4 processor with 3.0GHz clock speed

2018: Apple's MacBook Pro has clock speed of 2.7GHz

14 years later, not much has gain in raw processing power

Modern big data architects are required to "think outside of the box". Literally!

42 Years of Microprocessor Trend Data

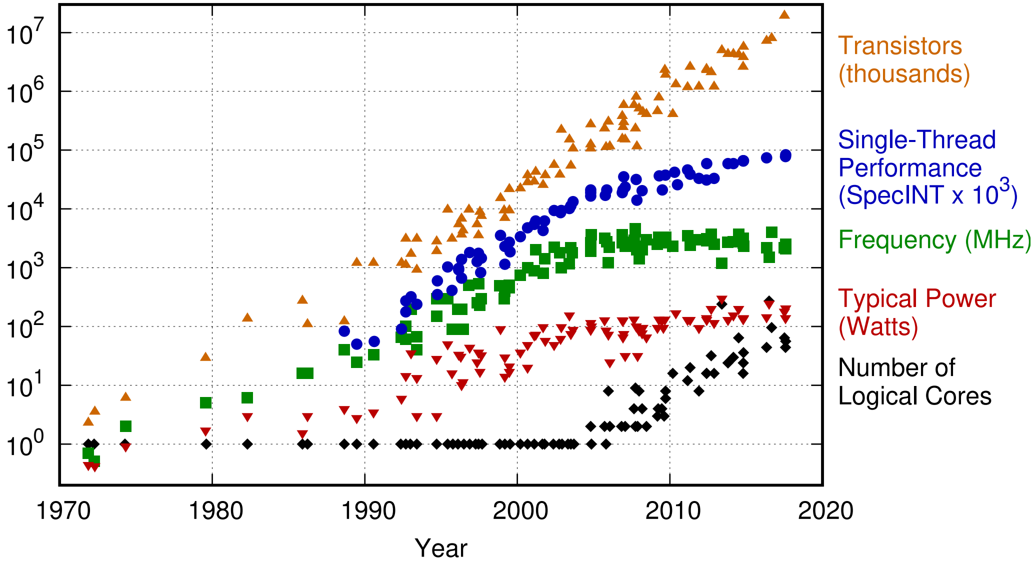

Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2017 by K. Rupp

Some Background Info

- **Agencies are historically focused on systematic capture and stewardship of data for observational Systems**
- **With large amount of observational and modeling data,**
  - The overall cost for data stewardship is expecting to rise significantly
  - Finding and downloading is becoming inefficient
- **Reality with large amount of observational and modeling data**
  - Downloading to local machine is becoming inefficient
  - Search has gotten a lot faster, but finding the relevant measurement has becoming a very time consuming process
  - Analyze decades of regional measurement is labor-intensive and costly
- **Increasing “big data” era is driving needs to**
  - Scale computational and data infrastructures
  - Support new methods for deriving scientific inferences and shift towards integrated data analytics
  - Apply computational and data science across the lifecycle
- **Scalable Data Management**
  - Capture well-architected and curated data repositories based on well-defined data/information architectures
  - Architecting automated pipelines for data capture
- **Scalable Data Analytics**
  - Access and integration of highly distributed, heterogeneous data
  - Novel statistical approaches for data integration and fusion
  - Computation applied at the data sources
  - Algorithms for identifying and extracting interesting features and patterns

National Aeronautics and Space Administration
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Traditional Method for Analyze Satellite Measurements

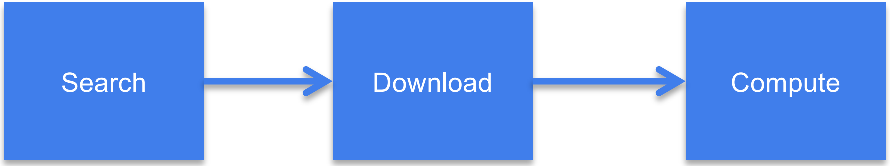

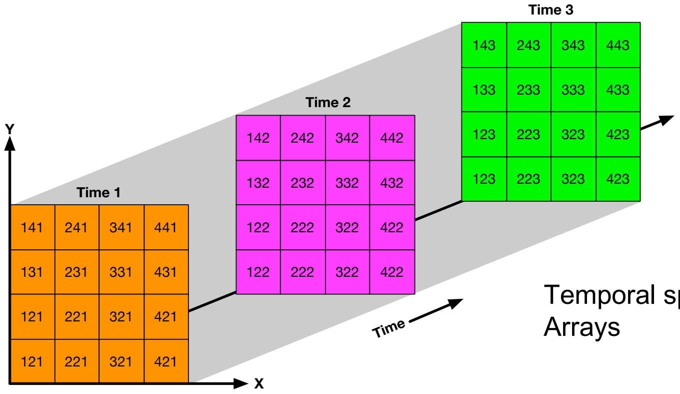

Temporal spatial Arrays

- Depending on the data volume (size and number of files)
- It could take many hours of download – (e.g. 10yr of observational data could yield thousands of files)
- It could take many hours of computation
- It requires expensive local computing resource (CPU + RAM + Storage)
- After result is produced, purge downloaded files

Observation

- Traditional methods for data analysis (time-series, distribution, climatology generation) can't scale to handle large volume, high-resolution data. They perform poorly
- Performance suffers when involve large files and/or large collection of files
- A high-performance data analysis solution must be free from file I/O bottleneck

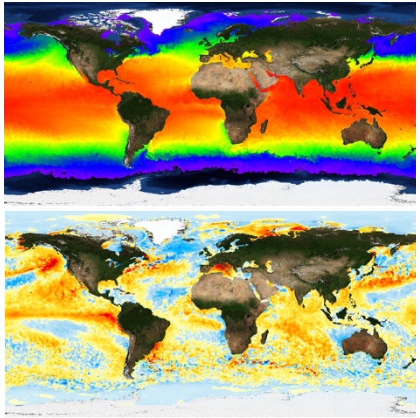

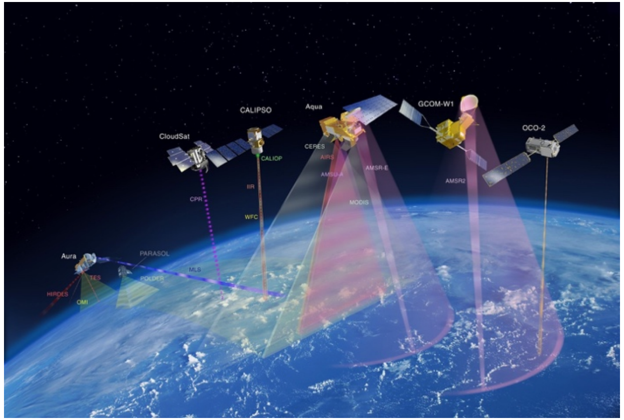

National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Building Climatological Services on the Cloud

T. Huang, et al, ICSE 2011, Honolulu, HI

- **2011 International Conference on Software Engineering (ICSE 2011)** - Software Engineering for Cloud Computing Workshop
- Publication of our first cloud-based analytics capability to analyze **SeaWinds on QuikSCAT Level 3 Sigma-O Polar-Sterographic Browse Maps of Antarctica, 25Km, 1999-2009**, 567,550,880 data points
- Developed tile-based analytics solution using Apache Hadoop and HBase

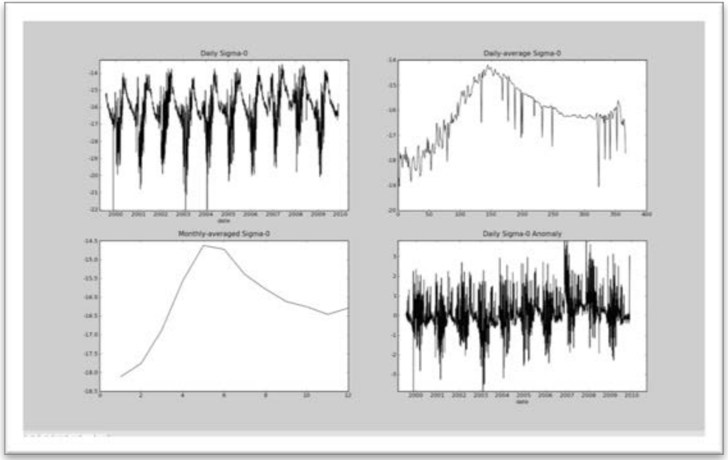

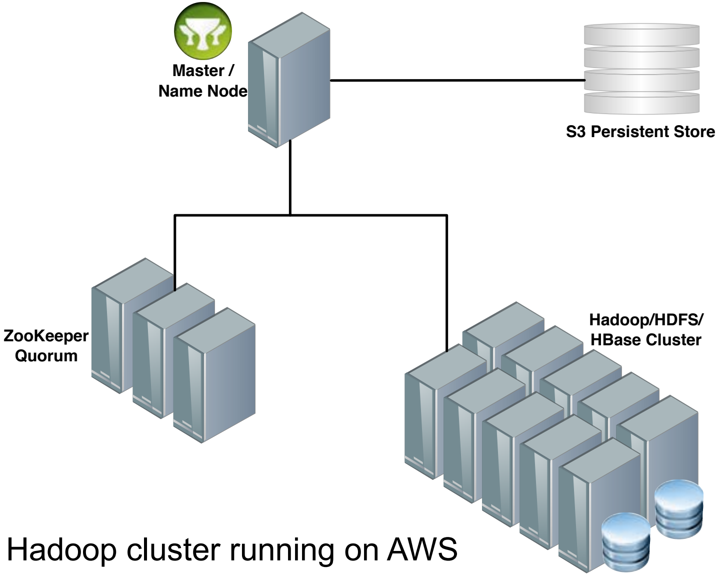

Hadoop cluster running on AWS

amazon web services

National Aeronautics and Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Scalable Data Analytic Solution

- **MapReduce**: A programming model for expressing distributed computations on massive amount of data and an execution framework for large-scale data processing on clusters of commodity servers. - J. Lin and C. Dyer, "*Data-Intensive Text Processing with MapReduce*"
  - **Map**: splits processing across cluster of machines in parallel, each is responsible for a record of data
  - **Reduce**: combines the results from Map processes
- In-memory parallel analytic solution using a new approach for handling science data to enable large-scale data analysis
  - Streaming architecture for horizontal scale data ingestion
  - Scales horizontally to handle massive amount of data in parallel
  - Provides high-performance geospatial and indexed search solution
  - Provides tiled data storage architecture to eliminate file I/O overhead
  - A growing collection of science analysis webservices

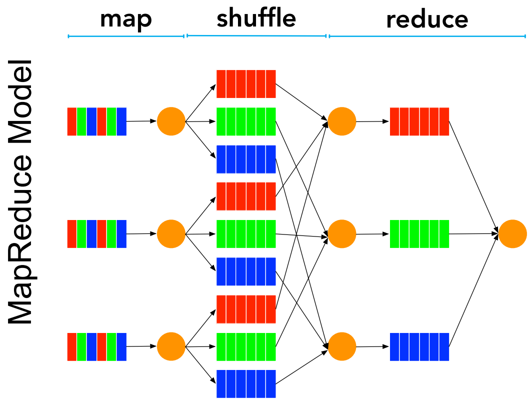

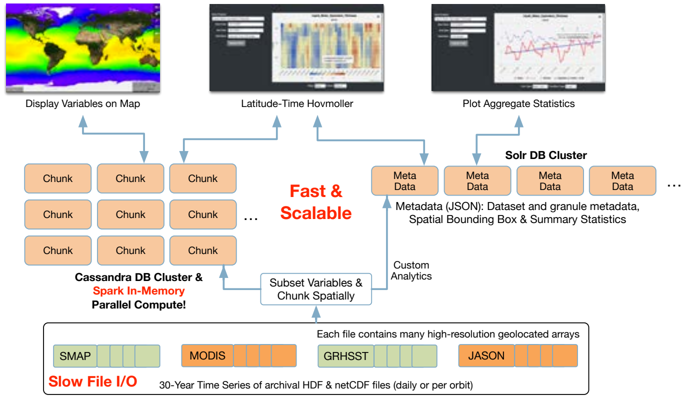

Apache NEXUS' Two-Database Architecture

NEXUS Performance: GIOVANNI vs. Custom Spark vs. AWS EMR

Dataset: MODIS AQUA Daily
Name: Aerosol Optical Depth 550 nm (Dark Target) (MYD08_D3v6)
File Count: 5106
Volume: 2.6GB
Time Coverage: July 4, 2002 – July 3, 2016

Giovanni: A web-based application for visualize, analyze, and access vast amounts of Earth science remote sensing data without having to download the data.

- Represents current state of data analysis technology, by processing one file at a time
- Backed by the popular NCO library. Highly optimized C/C++ library

AWS EMR: Amazon's provisioned MapReduce cluster

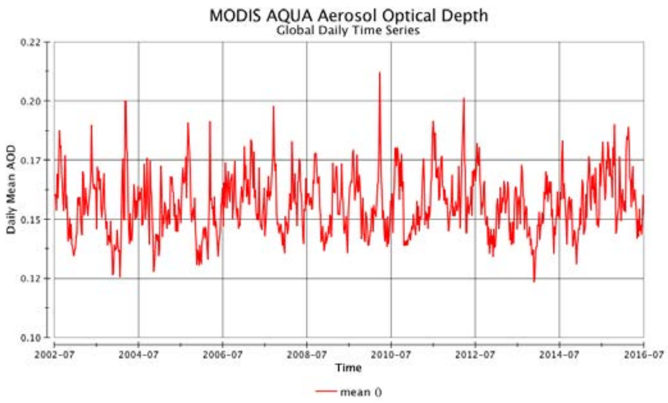

Giovanni: 20 min
NEXUS: 1.7 sec

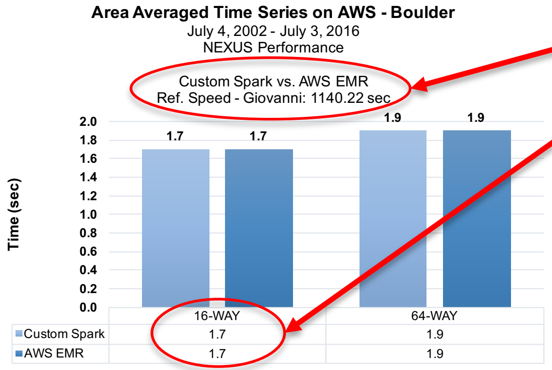

	16-WAY	64-WAY
Custom Spark	1.7	1.9
AWS EMR	1.7	1.9

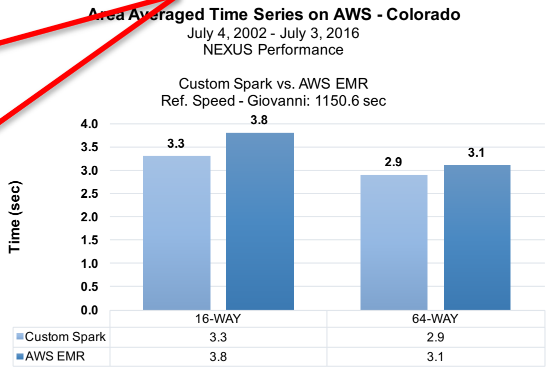

	16-WAY	64-WAY
Custom Spark	3.3	2.9
AWS EMR	3.8	3.1

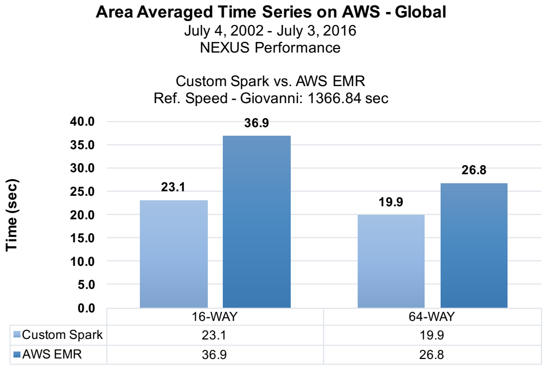

	16-WAY	64-WAY
Custom Spark	23.1	19.9
AWS EMR	36.9	26.8

Algorithm execution time. Excludes Giovanni's data scrubbing processing time

National Aeronautics and Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Integrated Science Data Analytics Platform

Creating SaaS and PaaS for Science Tools and Services

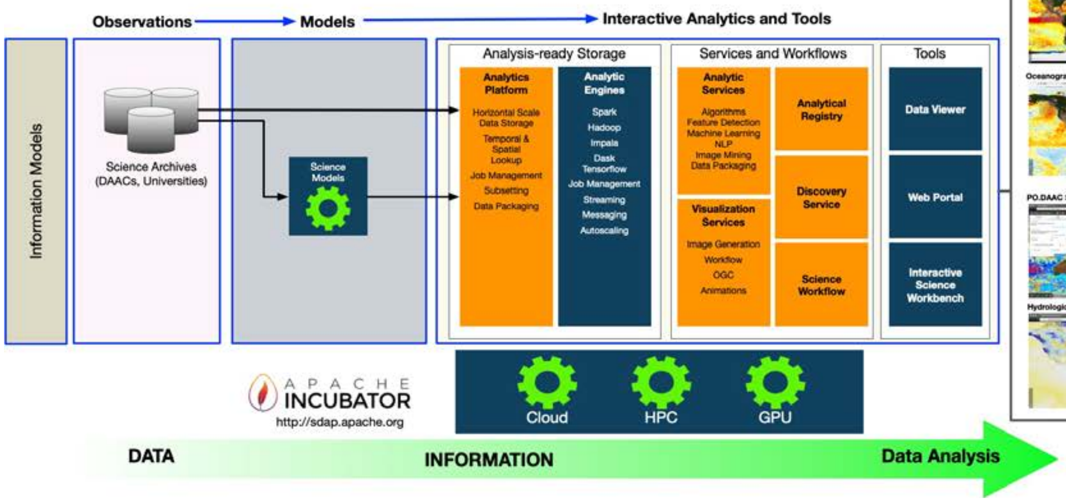

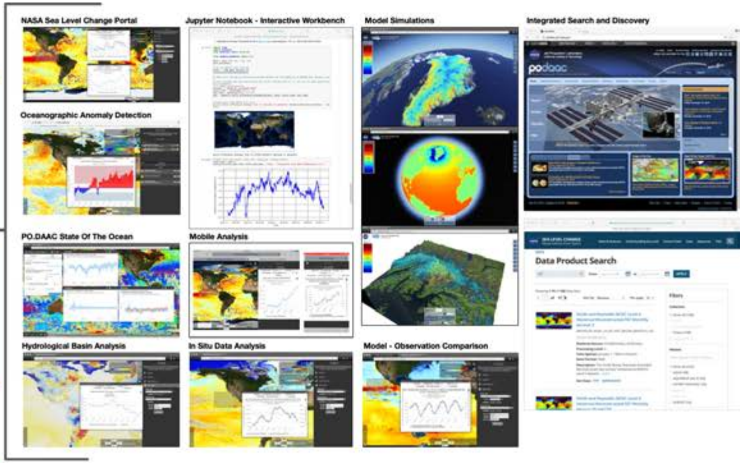

Systematic analysis of massive data

- **Integrated Science Data Analytics Platform**: an environment for conducting a science investigation
  - Enables the confluence of resources for that investigation
  - Tailored to the individual study area (physical ocean, sea level, etc.)
- Harmonizes data, tools and computational resources to permit the research community to focus on the investigation
- Scale computational and data infrastructures
- Shift towards integrated data analytics
- Algorithms for identifying and extracting interesting features and patterns

Analyze in situ and satellite observations

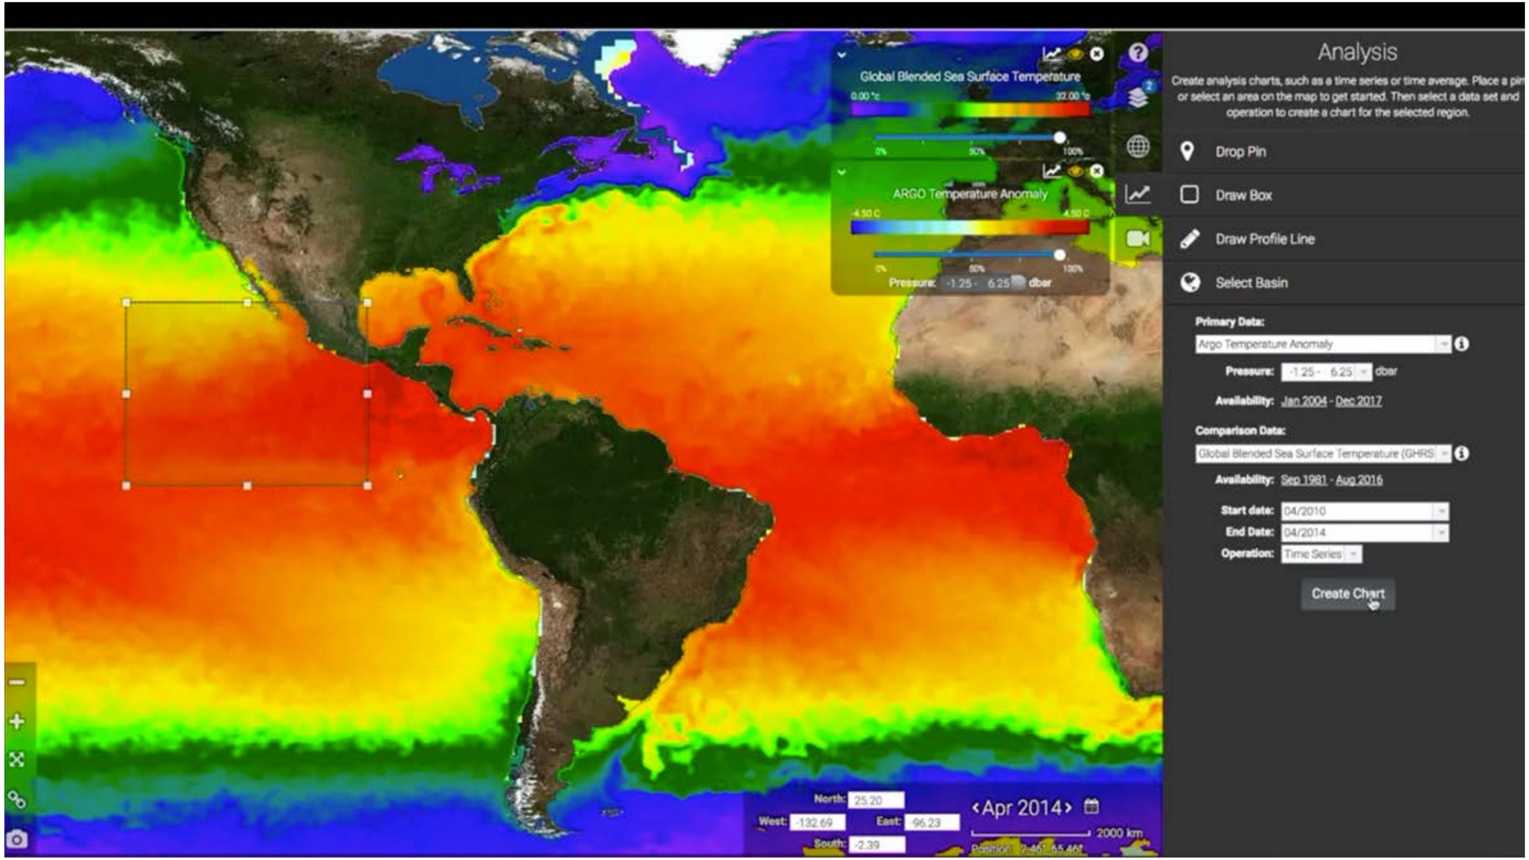

Mobile - https://sealevel.nasa.gov/data-analysis-tool/

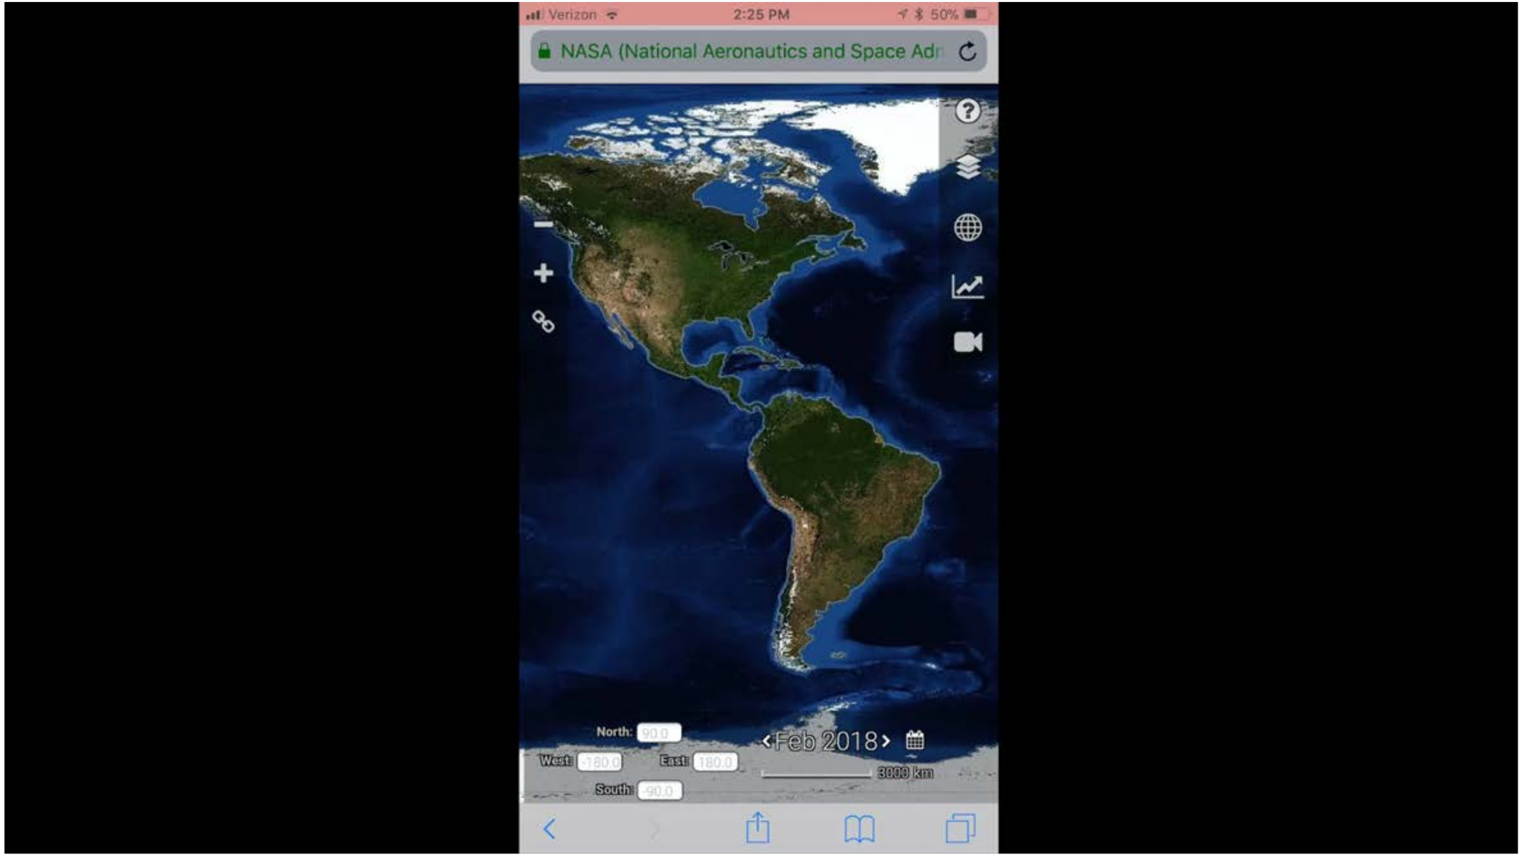

National Aeronautics and Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Lesson #1: Know The User's Real Needs

- **Work on improving communication - building bridge between IT and science**
  - **JPL's Data Science Working Group** is consists of technologists, project scientists, mission operations, etc.
  - Our science users tends get overwhelmed by tech jargons and cloud terminology
  - Learn to develop common language
- **Understand** how and for what purposes users obtain data and information
- **Describe** users' pain points and unmet needs for extracting, visualizing, comparing and analyzing science data
- **Identify** architectural approaches for tackling the real needs and identify opportunities for enhancing cross-disciplinary collaborative activities on the web portal.

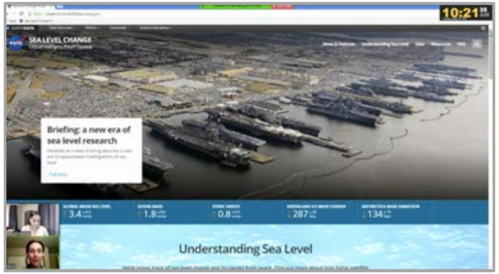

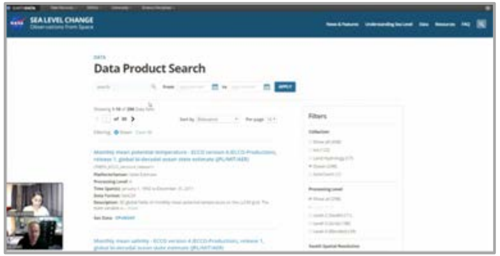

National Aeronautics and Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Hurricane Katrina Study

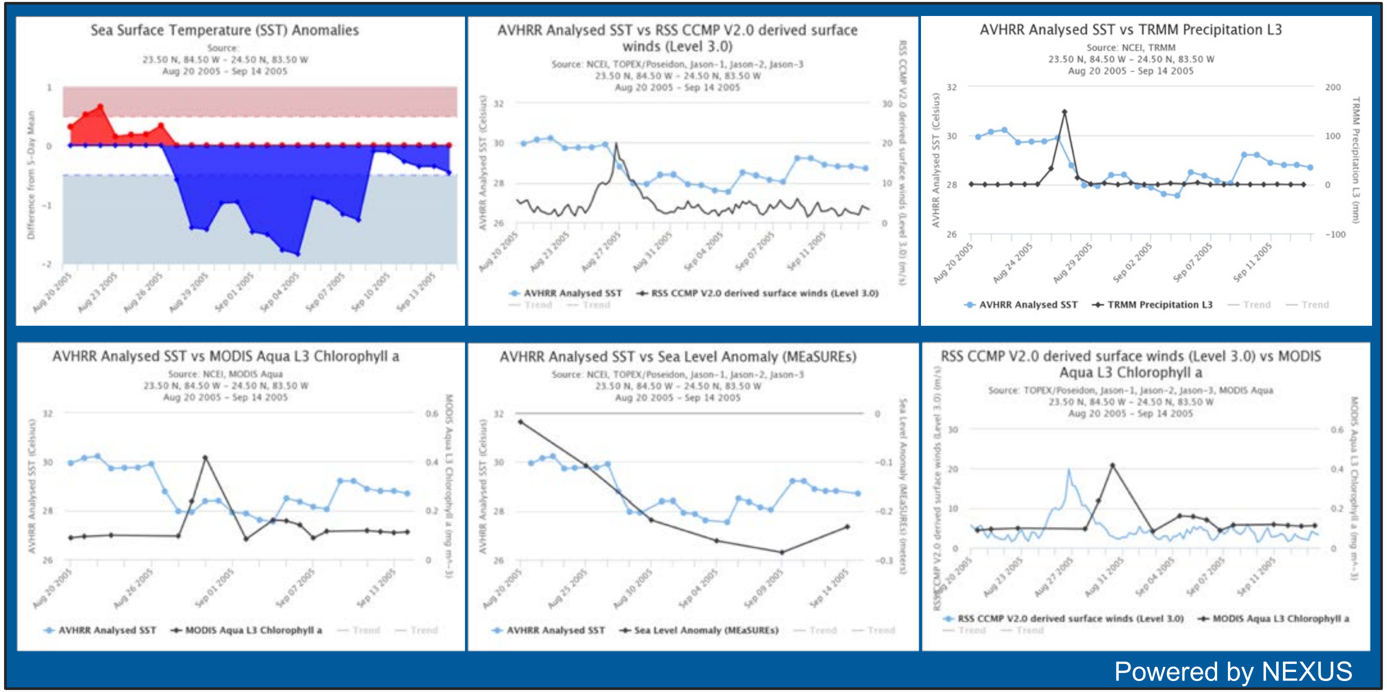

Hurricane Katrina passed to the southwest of Florida on Aug 27, 2005. The ocean response in a 1 x 1 deg region is captured by a number of satellites. The initial ocean response was an immediate cooling of the surface waters by 2 ℃ that lingers for several days. Following this was a short intense ocean chlorophyll bloom a few days later. The ocean may have been "preconditioned' by a cool core eddy and low sea surface height.

The SST drop is correlated to both wind and precipitation data. The Chl-A data is lagged by about 3 days to the other observations like SST, wind and precipitation.

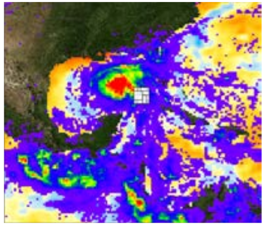

Hurricane Katrina TRMM overlay SST Anomaly

A study of a Hurricane Katrina–induced phytoplankton bloom using satellite observations and model simulations
Xiaoming Liu, Menghua Wang, and Wei Shi
JOURNAL OF GEOPHYSICAL RESEARCH, VOL. 114, C03023, doi:10.1029/2008JC004934, 2009

Lesson #2: Don't Forklift. Rearchitect.

- Forklift means moving current locally hosted systems to be hosted on the cloud
- Simply changing the hosting environment will not receive the benefits of the cloud. Actually, it could affect the performance and operating cost in a negative way
- Adding more computing node and/or cores does not always produce faster performance
- Rearchitect will be needed in order to fully benefit what cloud has to offer

Lesson #3: Research and Learn

- Fully understand the architecture of the system that you want to move to or develop
  - Workflow, resources, dependencies, security model
  - Storage type and size
  - Method of communication between processes and computing nodes
- Commercial Clouds provide many computing options at a pay-as-you-go model
  - **Computing**: types of CPU, GPU, number of cores, RAM, bus speed
  - **Storage**: SSD, spinning disk, attached, cross-mounted, I/O speed, write-once-read-many, frequency of access
  - **Network performance**: I/O speed, between computing nodes and storage, between on-premise and the cloud
  - **Pricing model**: on-demand, reserved, spot
  - **Provisioned solutions**: serverless, database, computing cluster, messaging, load-balancing and autoscaling, web caching, security
  - **Deployment process**: manual, VMs, containers, automate construction of infrastructure

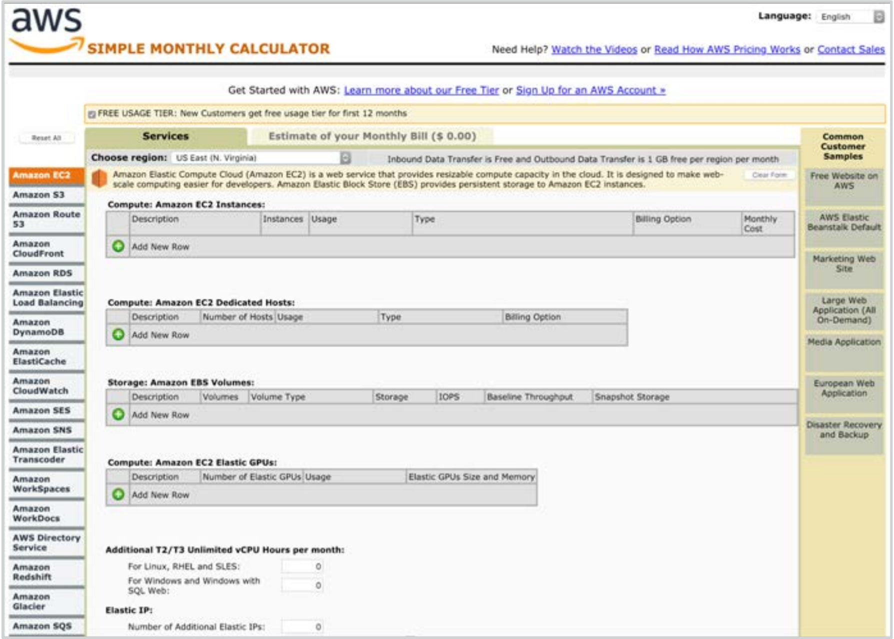

AWS Simple Monthly Calculator

Analytic Storage

Blocked
Storage
(EBS)

v.s.

Object
Storage
(S3)

$0.045/GB-month
$47,186/PB-month

$0.021/GB-month
$22,020/PB-month

Blocked storage (e.g. Amazon EBS), Attach to computing node. Generally faster

Object storage (e.g. Amazon S3), Independent storage service. Highly scalable

For a data center, not all data need to be served on fast storage, Object storage provides a better, scalable alternative

From Pilots to Operation

- From Hackers to Tinkers
- Prototypes and pilots for architectural validation
  - JPL's Data Science Program is funding internal pilots for this very reason
  - In our case, we needed a new approach for handling our large archive of NetCDFs and HDFs for analysis that reduces the data movement overhead
  - We changed how data is stored and architected our solution around parallel analytics
- Consider the end-to-end operation scenario from the beginning
  - How will operation team maintain this service?
  - Can we deploy patches without taking the entire system offline?
  - What is the total cost of ownership?
  - Cloud provides many options for computing to improve application speed, but can the project afford it?

NASA ACCESS 2017: ECCO-Cloud

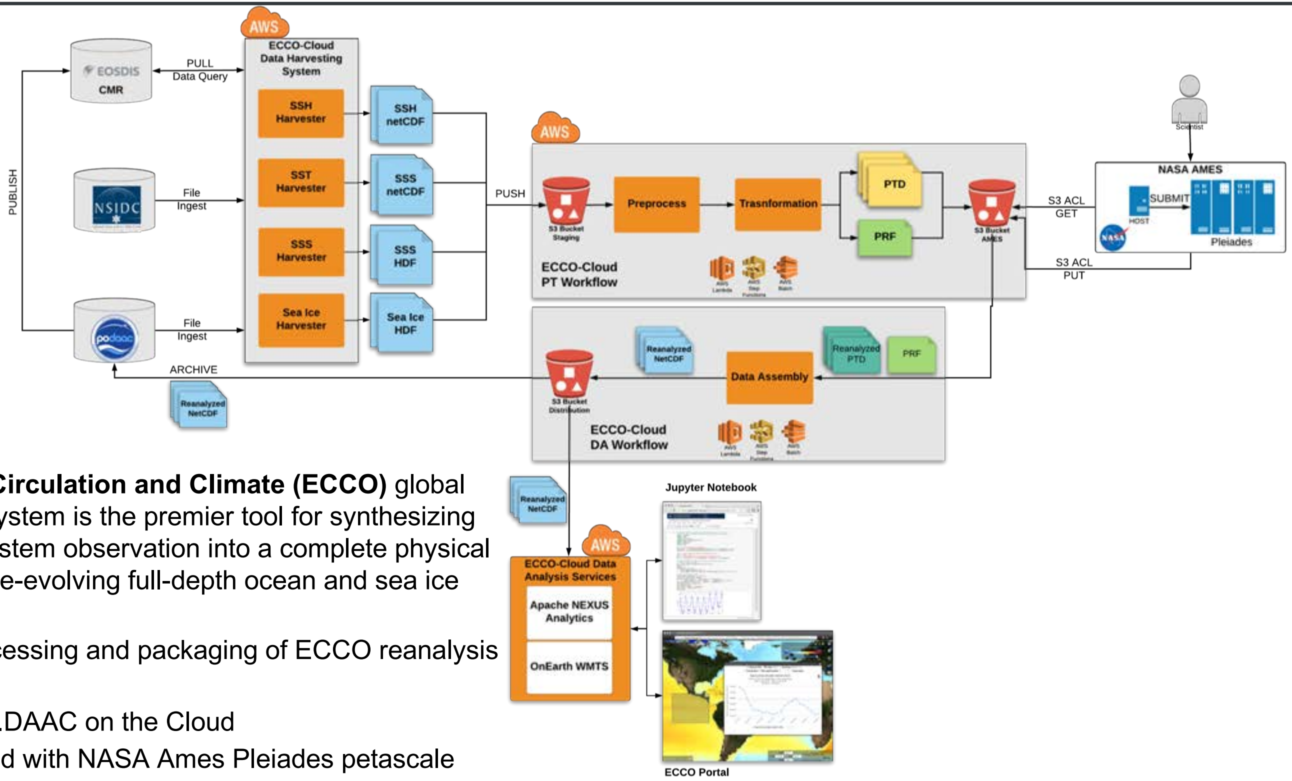

- **Estimating the Ocean Circulation and Climate (ECCO)** global ocean state estimation system is the premier tool for synthesizing NASA's diverse Earth system observation into a complete physical description of Earth's time-evolving full-depth ocean and sea ice system.
- Automate ingestion, processing and packaging of ECCO reanalysis products
- Automate delivery to PO.DAAC on the Cloud
- Integrating Amazon Cloud with NASA Ames Pleiades petascale supercomputer
- Establish ECCO Data Analysis Services and web portal for interactive visualization and analysis, and distribution using Apache SDAP

PI: Patrick Heimbach, University of Texas, Austin
Co-Is: Ian Fenty/JPL, Thomas Huang/JPL

National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Sea Level Analysis and ECCO

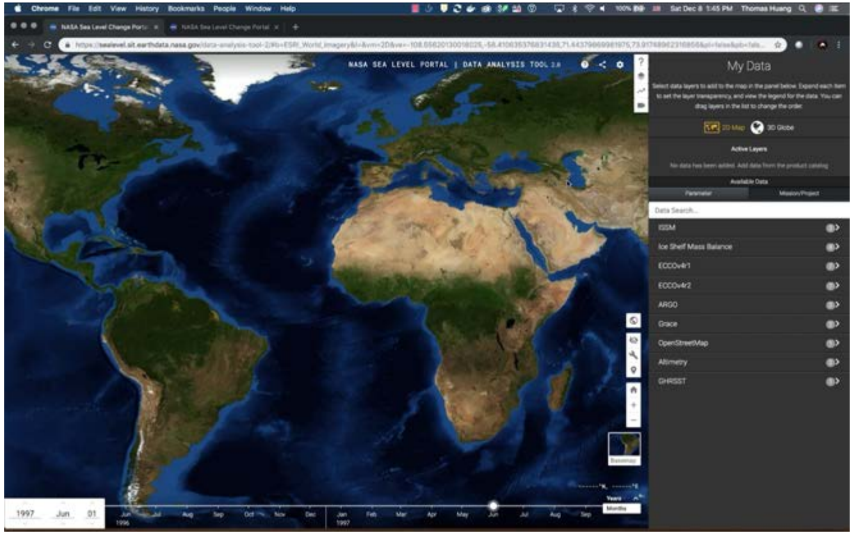

Dynamic colorbar and 3D visualization

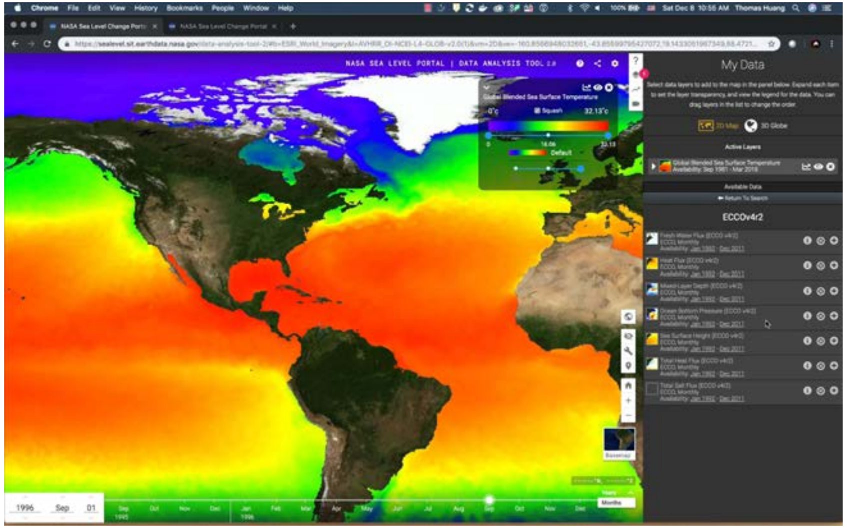

Overlay ECCO with satellite observation and 3D visualization

Evolve the Architecture

- **Several container-based deployment options**
  - Local on-premise cluster
  - Private Cloud
  - Amazon Web Service
- **Automate Data Ingestion with Image Generation**
  - Cluster based
  - Serverless (Amazon Lambda and Batch)
- **Data Store Options**
  - Apache Cassandra
  - ScyllaDB
  - Amazon Simple Storage Service (S3)
- **Resource Management Options**
  - Apache YARN
  - Apache MESOS
- **Analytic Engine Options**
  - Custom Apache Spark Cluster
  - Amazon Elastic MapReduce (EMR)
  - Amazon Athena (work-in-progress)

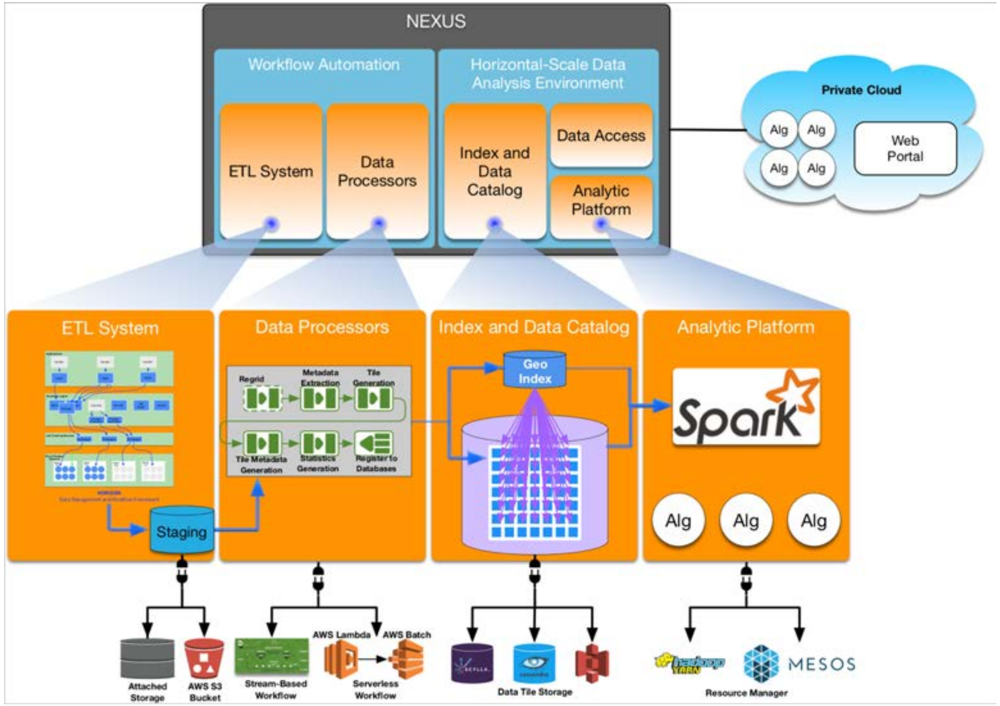

Apache NEXUS supports public/private Cloud and local cluster deployments

Lesson #4: Keep Service Interfaces Simple

- Encapsulation
  - Hiding what is under the hood simplifies integration
  - Improves portability
  - Keep vendor-specific logics hidden from service interfaces
- Programming language neutral web service interfaces, Applications can interact with the analytics platform using any programming languages (e.g. JavaScript, Python, Java, C/C++, IDL, MATLAB, etc.)
- Improve learning curve = Increase productivity
- Let our users focus on the science. Not technology!

National Aeronautics and Space Administration
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Meaningful Webservice URL Simplifies Integration

```
IDL> spawn, 'curl
"https://oceanworks.jpl.nasa.gov/timeSeriesSpark?spark=mesos,16,32&ds=AVHRR_OI_L4_GHRSS
T_NCEI&minLat=45&minLon=-150&maxLat=60&maxLon=-120&startTime=2008-09-
01T00:00:00Z&endTime=2015-10-01T23:59:59Z" -o json_dump.txt'
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload   Total   Spent    Left  Speed

  0     0    0     0    0     0      0      0 --:--:-- --:--:-- --:--:--     0
  0     0    0     0    0     0      0      0 --:--:--  0:00:01 --:--:--     0
  0     0    0     0    0     0      0      0 --:--:--  0:00:02 --:--:--     0
  0     0    0     0    0     0      0      0 --:--:--  0:00:03 --:--:--     0
  0  353k    0     0    0     0      0      0 --:--:--  0:00:03 --:--:--     0
 72  353k   72  256k    0     0  52705      0  0:00:06  0:00:04  0:00:02 52702
100  353k  100  353k    0     0  69303      0  0:00:05  0:00:05 --:--:-- 98883
IDL>
IDL> result = JSON_PARSE( 'json_dump.txt', /toarray, /tostruct)
IDL> help, result
** Structure <1a2749c8>, 3 tags, length=62320, data length=62320, refs=1:
   STATS           STRING    '!NULL'
   META            STRUCT    -> <Anonymous> Array[1]
   DATA            STRUCT    -> <Anonymous> Array[778]
IDL>
IDL> plot(result.data.time, result.data.mean, title='GHRSST L4 AVHRR_OI SST.  Sep 2008
- Oct 2015.   US West Coast Blob Area')
PLOT <29457>
```

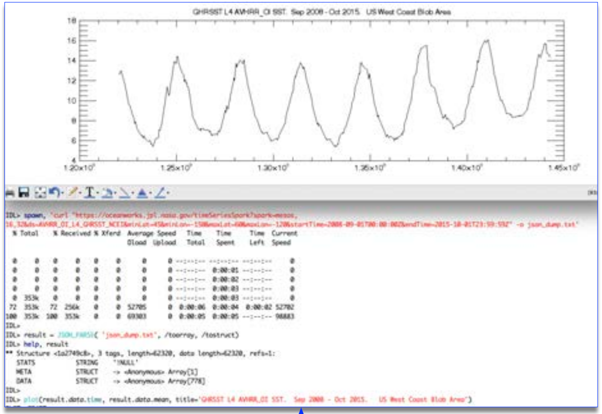

Credit: Ed Armstrong

Jun. 05, 2018

National Aeronautics and Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Distributed Analytics

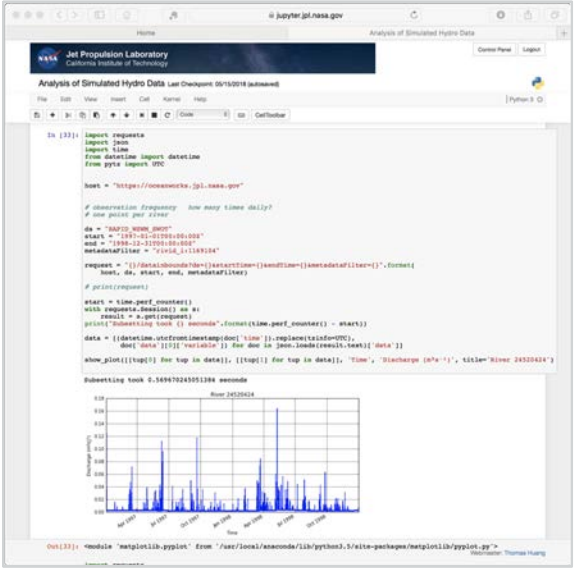

Retrieval of a single river time series

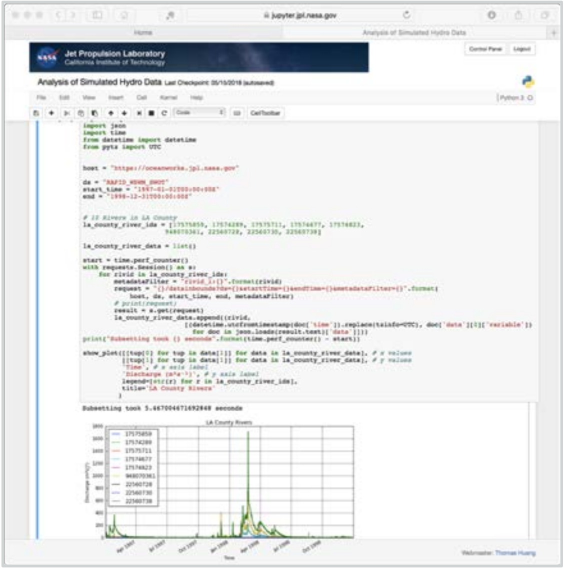

Retrieval of time series from 9 rivers

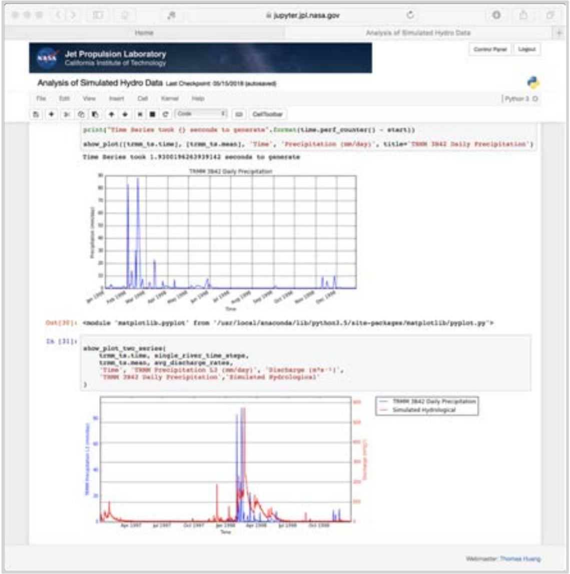

Time series coordination between TRMM and river

- Accessing analytic services hosted at JPL and on Amazon through simple interfaces
- Simulated hydrology data in preparation for SWOT hydrology
- **River data**: ~3.6 billion data points. 3-hour sample rate. Consists of measurements from ~600,000 rivers
- **TRMM data**: 17 years, .25deg, 1.5 billion data points
- Sub-second retrieval of river measurements
- On-the-fly computation of time series and generate coordination plot

National Aeronautics and Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Support Algorithm Development

- Scientific algorithm developers are power users
- They know how to work with the data, but they probably don't want to know the internals of the big data architecture
- We created service interfaces for users to inject custom algorithmic codes into our cloud services
- Architecture for sharing large computing engines to reduce operating cost

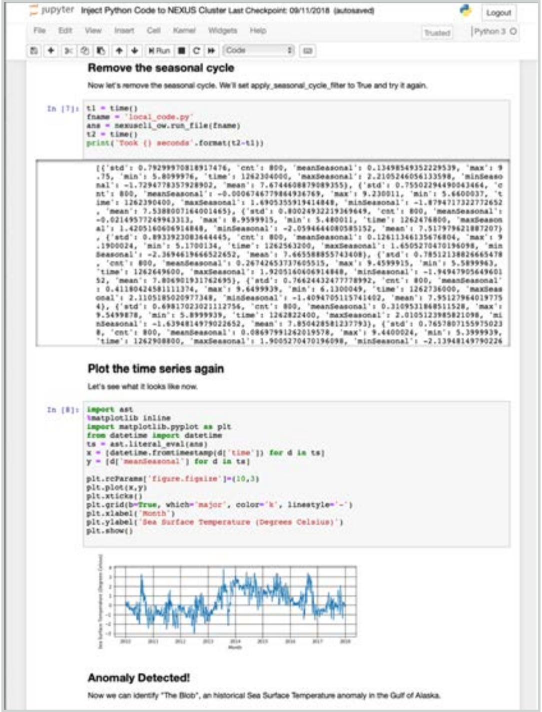

Injecting custom anomaly detection algorithm into our cloud-based analytic engine – J. Jacob, 2018

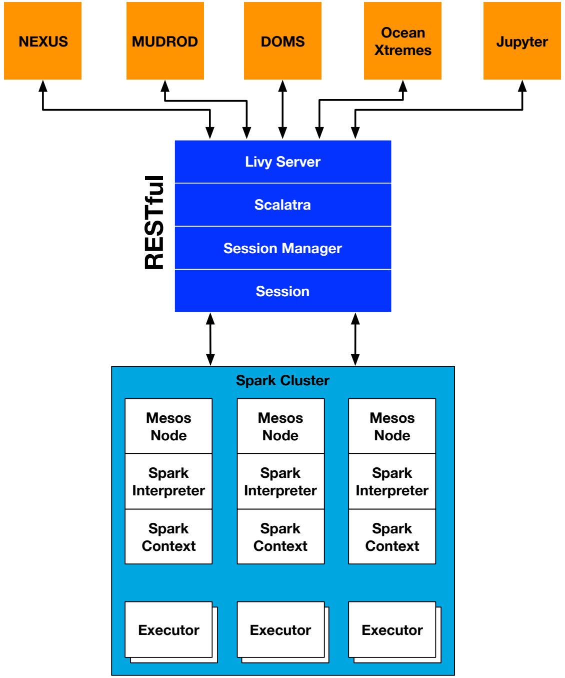

Example architecture for integrating different analytic services to share a common sharing Apache Spark engine = Cost Saving

In Situ to Satellite Matchup

- Distributed Oceanographic Matchup Service (DOMS)
- Typically data matching is done using one-off programs developed at multiple institutions
- A primary advantage of DOMS is the reduction in duplicate development and man hours required to match satellite/in situ data
    - Removes the need for satellite and in situ data to be collocated on a single server
    - Systematically recreate matchups if either in situ or satellite products are re-processed (new versions), i.e., matchup archives are always up-to-date.
- In situ data nodes at JPL, NCAR, and FSU operational.
- Provides data querying, subset creation, match-up services, and file delivery operational.
- Plugin architecture for in situ data source

```
url = 'https://doms.jpl.nasa.gov/match_spark?'
url += 'primary=' + primary
url += '&matchup=' + secondary
url += '&startTime='+ startTime
url += '&endTime=' + endTime
url += '&tt=' + timeTolerance
url += '&rt=' + radius
url += '&b=' + bbox
url += '&platforms=' + platforms
url += '&parameter=' + parameter
url += '&matchOne=true'
url += '&depthMin=' + depthMin
url += '&depthMax=' + depthMax

print (url)
start = time.time()
ts = json.loads(str(requests.get(url).text))
spent = time.time() - start
print ("It took: " + str(spent) + " sec")

satellite = []
in_situ = []
for data in ts ['data']:
    for matches in data ['matches']:
        satellite.append(data['sea_water_temperature'])
        in_situ.append(matches['sea_water_temperature'])

show_plot(in_situ, satellite, secondary+' (c)', primary+' (c)')
```

```
https://doms.jpl.nasa.gov/match_spark?primary=AVHRR_OI_L4_GHRSST_NCEI&matchup=spurs&startTime=2013-10-01T00:00:00Z&endTime=2013-10-30T23:59:59Z&tt=86400&rt=10000.0&b=-30,15,-45,30&platforms=1,2,3,4,5,6,7,8,9&parameter=sst&matchOne=true&depthMin=0&depthMax=5
It took: 2.4485230445861816 sec
```

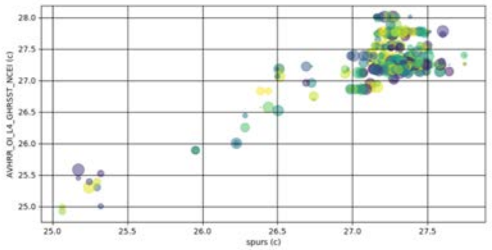

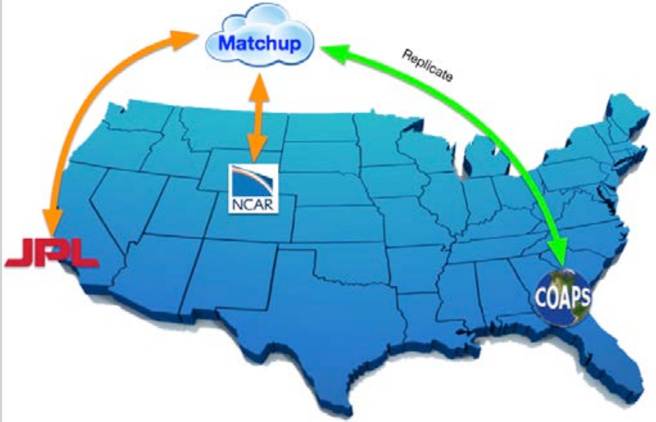

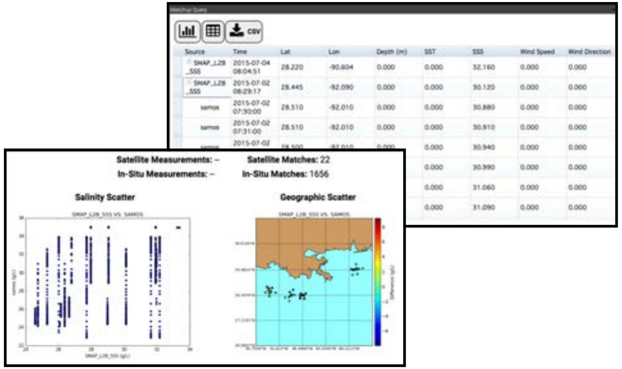

National Aeronautics and Space Administration
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Architecture for Distributed Data System and Analysis

Committee on Earth Observation Satellites (CEOS)

- **CEOS Ocean Variables Enabling Research and Applications for GEO (COVERAGE) Initiative**
- Seeks to provide **improved access** to **multi-agency ocean remote sensing data** that are **better integrated with in-situ and biological observations**, in support of **oceanographic and decision support applications** for societal benefit.
- A community-support open specification with common taxonomies, information model, and API (maybe security)
- Putting value-added services next to the data to eliminate unnecessary data movement
- Avoid data replication. Reduce unnecessary data movement and egress charges
- Public accessible RESTful analytic APIs where computation is next to the data
- Analytic engine infused and managed by the data centers perhaps on the Cloud
- Researchers can perform multi-variable analysis using any web-enabled devices without having to download files

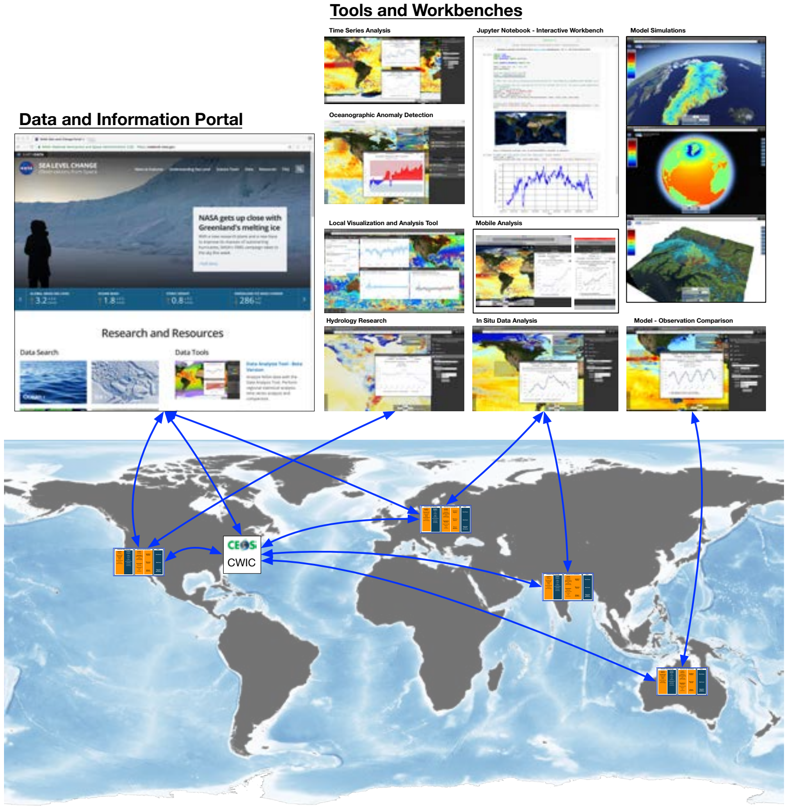

Lesson #5: Streamline Deployment

- Cloud Deployment is nontrivial
- Infrastructure Definition
  - Various machine instances
  - Storage and buckets
- Software Deployment.. manually
  - Build
  - Package
  - Install
  - Configure
  - Shell login (security issues)
- Best Practice: Deployment Automation
  - Script Infrastructure Definition (e.g. Amazon CloudFormation)
  - Container-based Deployment (e.g. Amazon ECS and DockerHub)

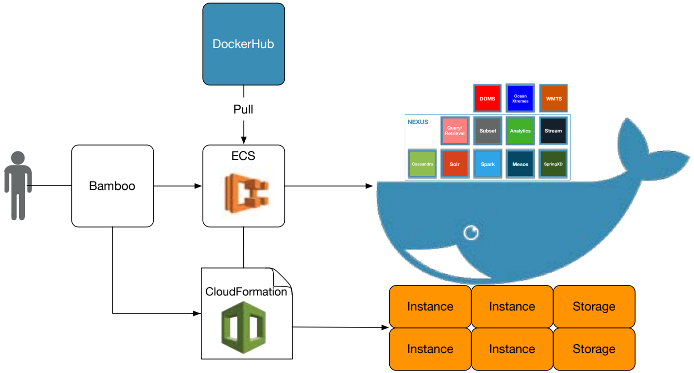

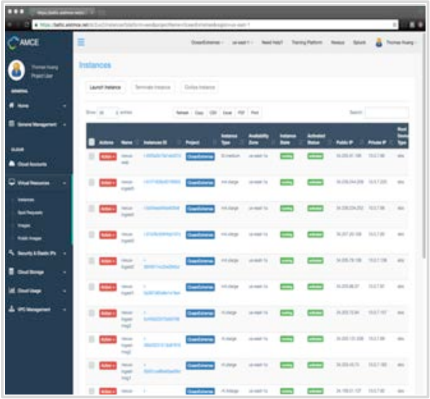

AIST Managed Cloud Environment

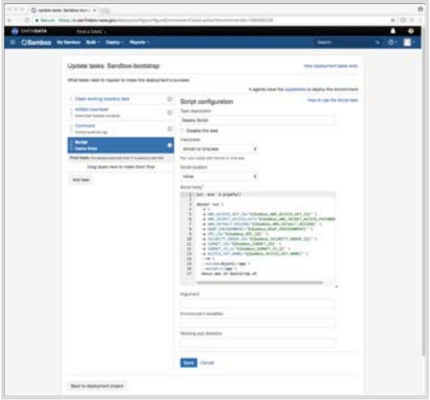

ESDIS NGAP

Lesson #6: Building Community-Driven Open Source Solution

- Develop in the open, so every data provider can infuse the same software stack next to their data
- Establish or leverage an existing governance policy
- Community accessible issue tracking and documentations
- Community validation
- Evolve the technology through community contributions
- Share recipes and lessons learned
- Remember open source != less secure. Some open source technologies, Linux, Apache Webserver,, GNU, etc., have already been adopted by enterprises for years
- Host webinars, hands-on cloud analytics workshops and hackathons



Big Data Analytics and Cloud Computing Workshop, 2017 ESIP Summer Meeting, Bloomington, IN

National Aeronautics and Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Free and Open Open Source Software (FOSS)

- October 2017, the OceanWorks project released all of its source code to Apache Software Foundation and established the **Science Data Analytics Platform (SDAP)** in the **Apache Incubator**
- Technology sharing through Free and Open Source Software (FOSS)
- Why? Further technology evolution that is restricted by projects / missions
- It is more than GitHub
  - Quarterly reporting
  - Reports are open for community review by over 6000 committers
  - SDAP has a group of appointed international Mentors: Jörn Rottmann, Raphael Bircher, and Suneel Marthi
- OceanWorks is now being developed in the open
  - For local cluster and cloud computing platform
  - Fully containerized using Docker (multiple containers)
  - Infrastructure orchestration using Amazon CloudFormation
  - Analyzing satellite and model data
  - In situ data analysis and colocation with satellite measurements
  - Fast data subsetting
  - Data services integration architecture
  - OpenSearch and dynamic metadata translation
  - Mining of user interactions and data to enable discovery and recommendations
  - Streamline deployment through container technology

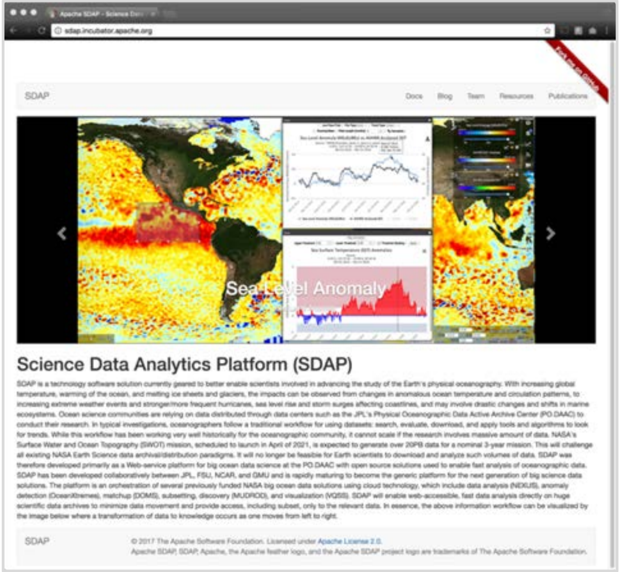

http://sdap.apache.org

APACHE INCUBATOR

National Aeronautics and Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Lesson #7: Deliver to Production

- The gap between visionary to pragmatists is significant. It must be the primary focus of any long-term high-tech marketing plan – Geoffrey Moore
- Become an expert in the production environment and devote resources in creating automations
- Give project engineering team early access to the PaaS
- Deliver all technical documents and work with project system engineering
- Provide user-focused trainings

NASA's Physical Oceanography Distributed Active Archive Center (PO.DAAC)

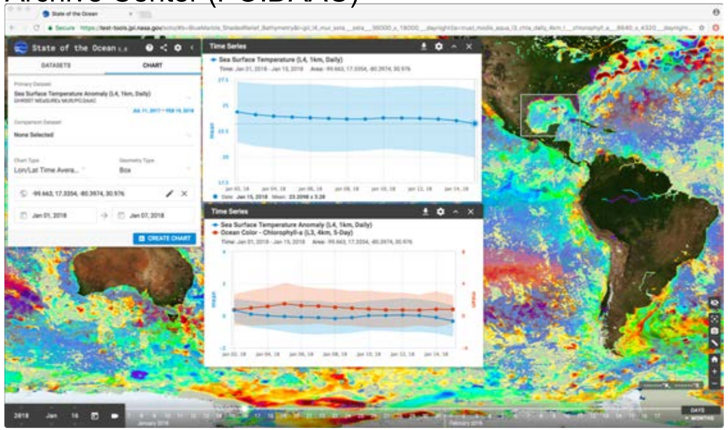

NASA Sea Level Change Team

In Summary

- Our approach to big data analytics is disruptive to Data Huggers and Server Huggers
- Our method of development is disruptive to Software Huggers
- Our motivation is to develop technology solution to fully leverage the cloud to deliver an answer to our users as quickly as funding allow
- Data linage and access to the original data is important, but we must understand the motivation for download.
- Can we develop and provide the solutions to meet some of these needs without having to download?
- Disruptive Innovations are products that require us to change our current mode of behavior or to modify other products and services – Geoffrey Moore
- We will see many disruptive innovations for tackling our big data analytics challenges

https://topspeeddata.com/spinning-top-blog/2018/3/30/dont-be-a-server-hugger

National Aeronautics and Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Big thanks to our sponsors

- NASA ESTO/AIST Program
- NASA ACCESS Program
- NASA Sea Level Rise Program
- NASA ESDIS Project
- JPL Data Science Program

Thomas Huang
Thomas.huang@jpl.nasa.gov
Jet Propulsion Laboratory
California Institute of Technology