

7th Annual
Earth System Grid Federation
December 2017



is-enes



Face-to-Face Conference

A global consortium of government agencies, educational institutions, and companies dedicated to delivering robust distributed data, computing libraries, applications, and computational platforms for the novel examination of extreme-scale scientific data.

STATE OF THE EARTH SYSTEM GRID FEDERATION (ESGF) 2017

ESGF EXECUTIVE COMMITTEE

DEAN WILLIAMS (DOE/LLNL, CHAIR)

MICHAEL LAUTENSCHLÄGER (DKRZ, CO-CHAIR)

SEBASTIEN DENVIL (IPSL)

ROBERT FERRARO (NASA/JPL)

DANIEL DUFFY (NASA/GSFC)

V. BALAJI (NOAA/GFDL)

BEN EVANS (NCI)

TOM LANDRY (CRIM)

LUCA CINQUINI (NASA/JPL)

ESGF: WINNER OF 2017 R&D 100 AWARD

- ▶ ESGF WAS ONE OF THE WINNERS OF 2017 R&D 100 AWARDS!
- ▶ EVERY YEAR, THIS PRESTIGIOUS CONFERENCE RECOGNIZES THE TOP 100 MOST INNOVATIVE PRODUCTS IN SOFTWARE, SCIENCE & TECHNOLOGY
- ▶ ESGF WON IN THE COMPUTER SCIENCE CATEGORY FOR "OUTSTANDING SERVICE TO THE COMMUNITY"



ESGF F2F 2017 PART 1

PRESENT STATE OF THE EARTH SYSTEM GRID FEDERATION

ESGF OVERVIEW

- ▶ ESGF IS AN INTERNATIONAL COLLABORATION OF CLIMATE CENTERS WORKING TOGETHER TO MANAGE AND PROVIDE ACCESS TO CLIMATE DATA - MODELS AND OBSERVATIONS
- ▶ STARTED MORE THAN A DECADE AGO, NOW THE WORLD PREMIER TECHNOLOGY INFRASTRUCTURE IN SUPPORT OF CLIMATE SCIENCE
- ▶ SPANNING SEVERAL TENS OF INSTITUTIONS IN EUROPE, NORTH AMERICA, AUSTRALIA, ASIA, WITH FUNDING FROM MULTIPLE AGENCIES



is-enes
INFRASTRUCTURE FOR THE EUROPEAN NETWORK
FOR EARTH SYSTEM MODELLING



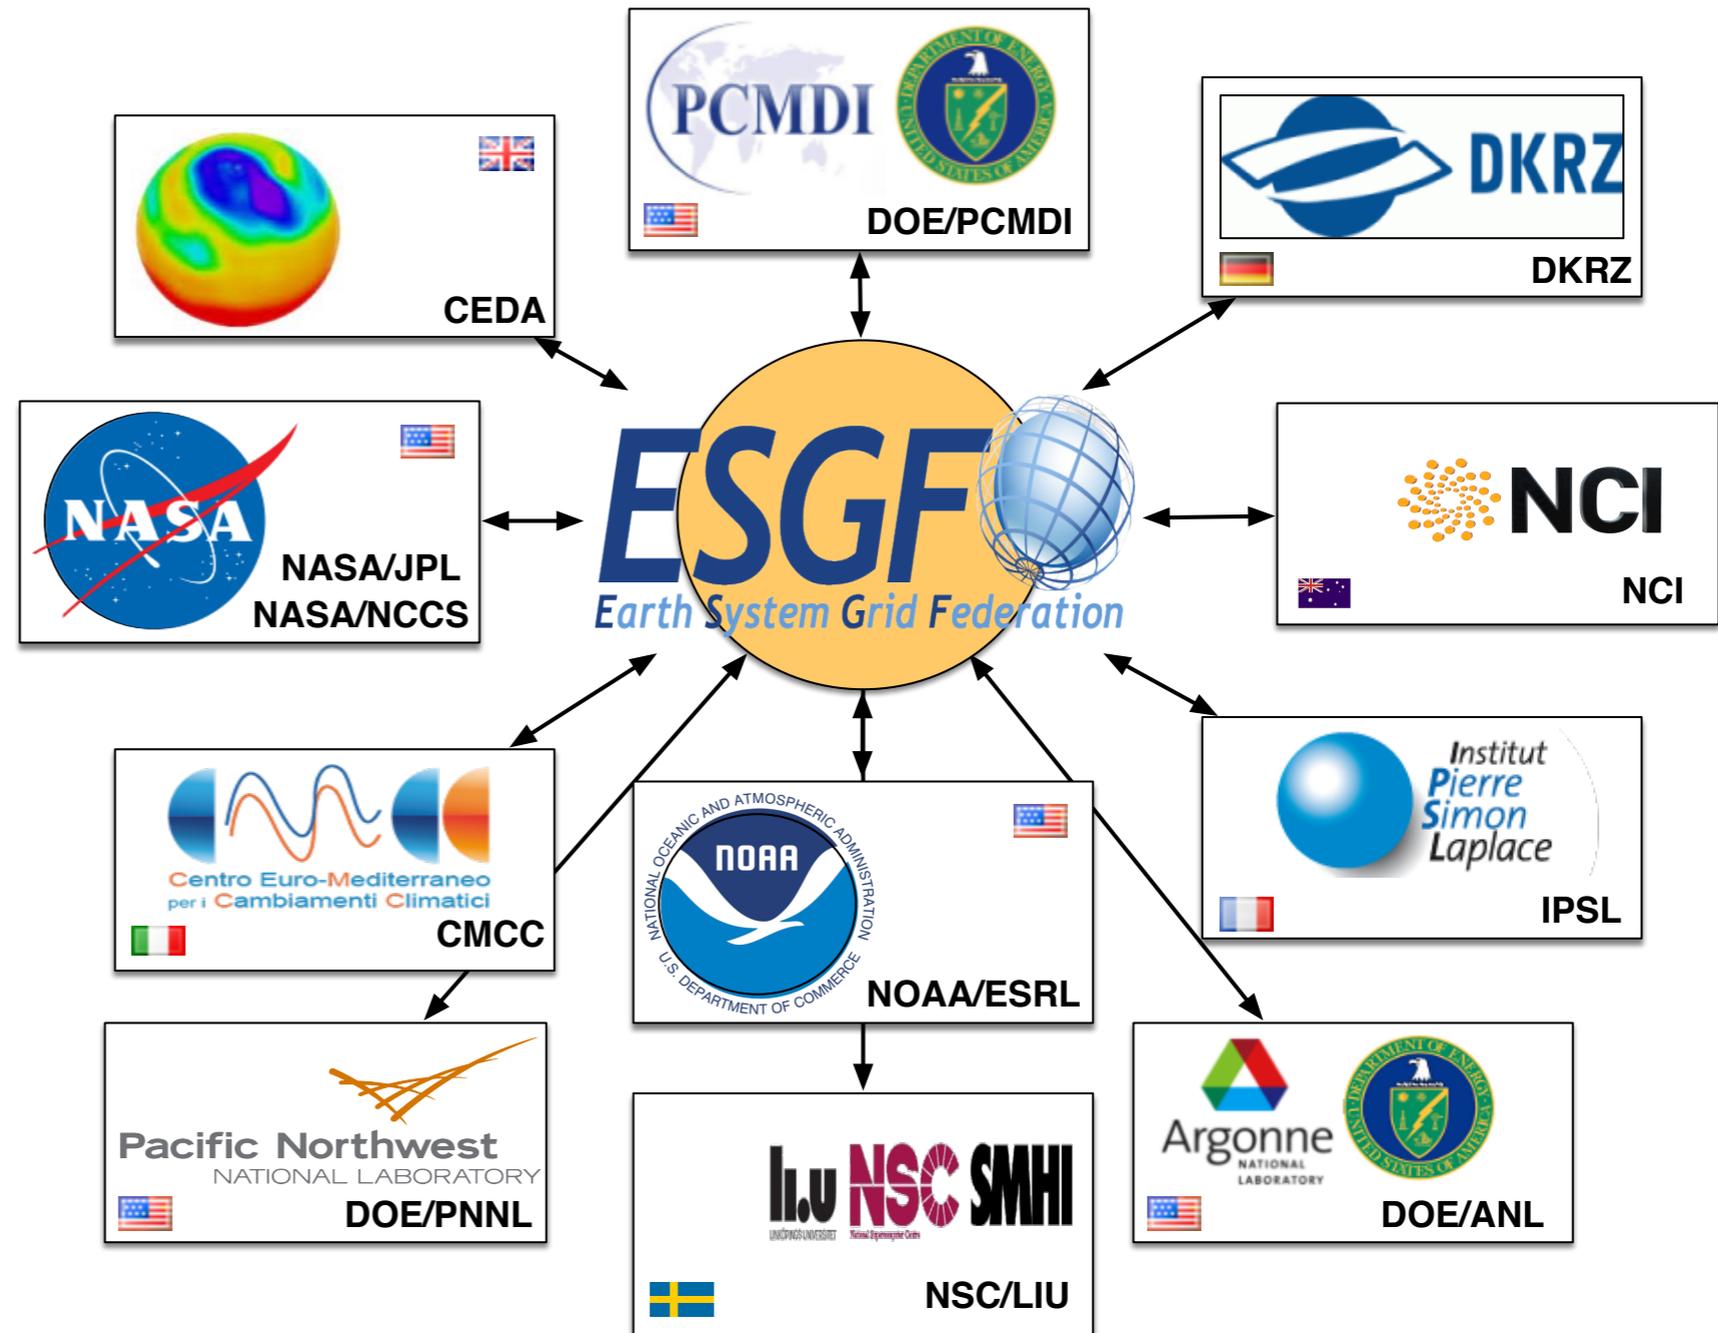
canarie



ESGF FEDERATION

- ▶ The real value of ESGF comes from the fact that we are a federation of sites: by using ESGF, users have access to data, metadata and services that are provided by multiple sites distributed around the world
- ▶ Interoperability across sites is based on:

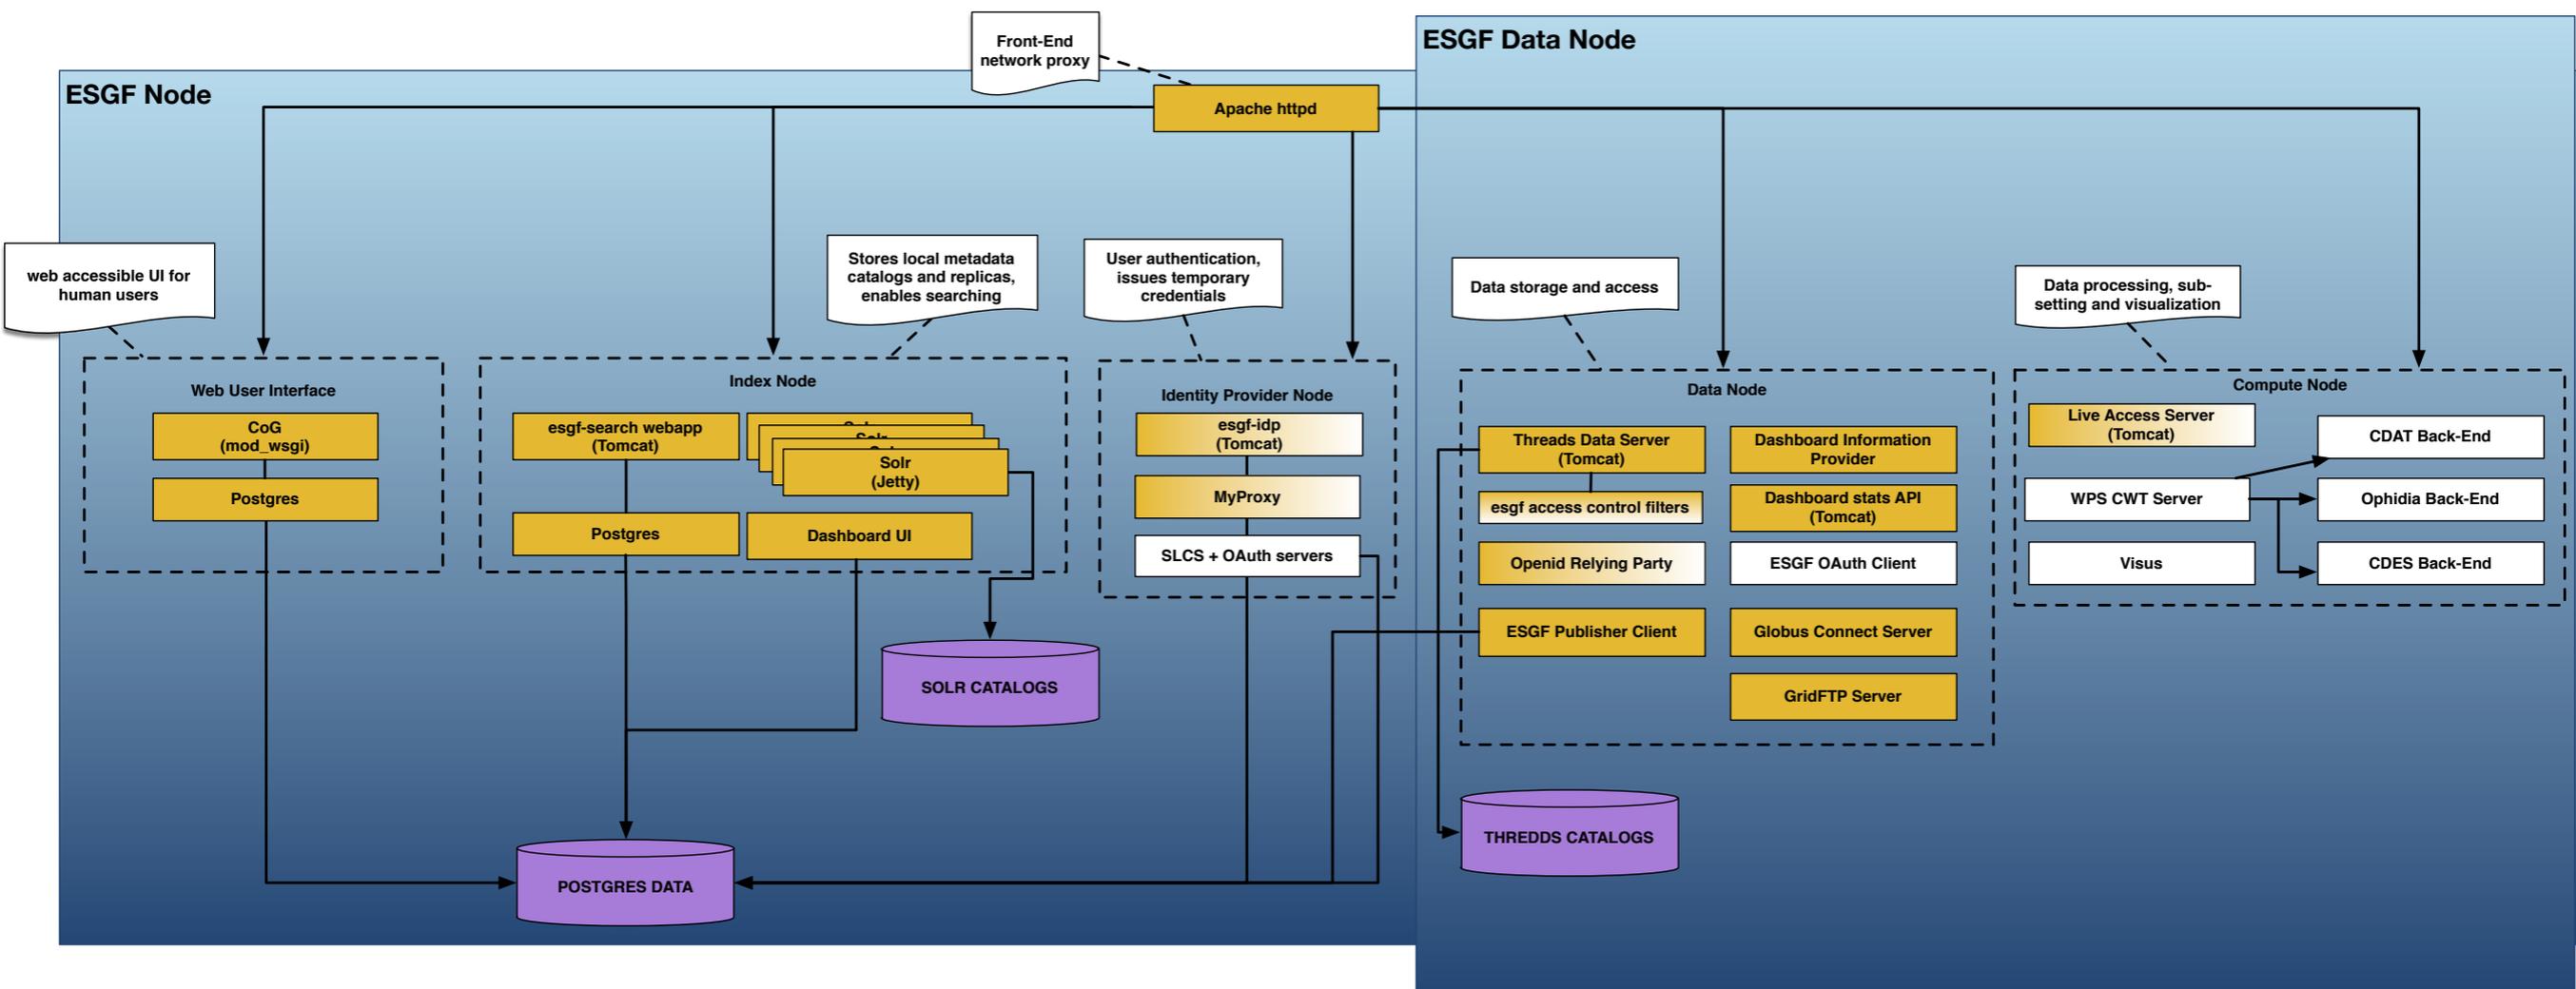
- ▶ Sites operate services that are based on standard protocols and APIs (SSL, OpenID, WPS, ...)
- ▶ Services are configured to know about their peers (e.g.: search, IdPs, CoGs, ...)
- ▶ Sites establish trust relationships through SSL certificates



- ▶ 31 data nodes, 11 index nodes
- ▶ 23 data projects, 198,298 datasets, 3,309,528 files (non-replica, latest version)
- ▶ 147 CoG projects, 15,571 users across all CoG sites

ESGF SOFTWARE STACK

- ▶ Integration of many applications and services, developed by ESGF or from Open Source
- ▶ Logically grouped in areas of functionality: User Interface, Index Node, Identity Provider, Data Node, Compute Node
- ▶ Based on standard protocols and APIs for interoperability
- ▶ Currently supporting 2 deployment architectures: complete ESGF Node, Data Node only
- ▶ Stack is evolving: currently replacing some deprecated components and adding new services



ESGF WORKING GROUPS

- ▶ ESGF software development is organized by working groups
- ▶ # of groups recently reduced to remove duplication of functionality and enable tighter collaboration
- ▶ Working groups report and discussion will represent the core of this conference
- ▶ Each working group is asked to provide a final oral and written report including:
 - ▶ progress to date
 - ▶ future roadmap

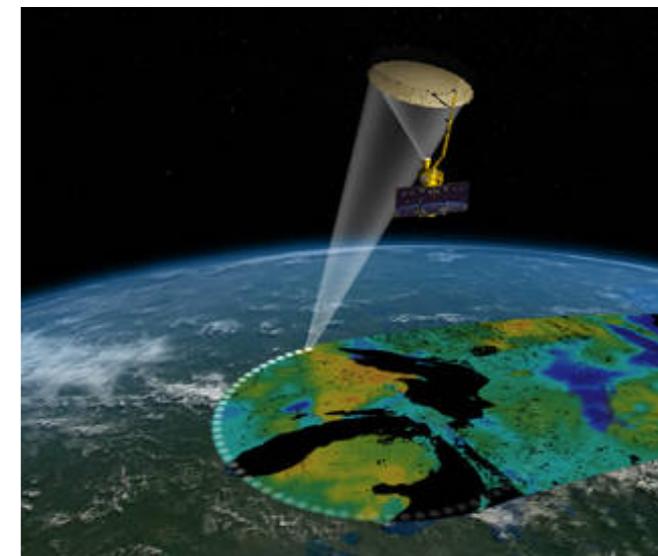
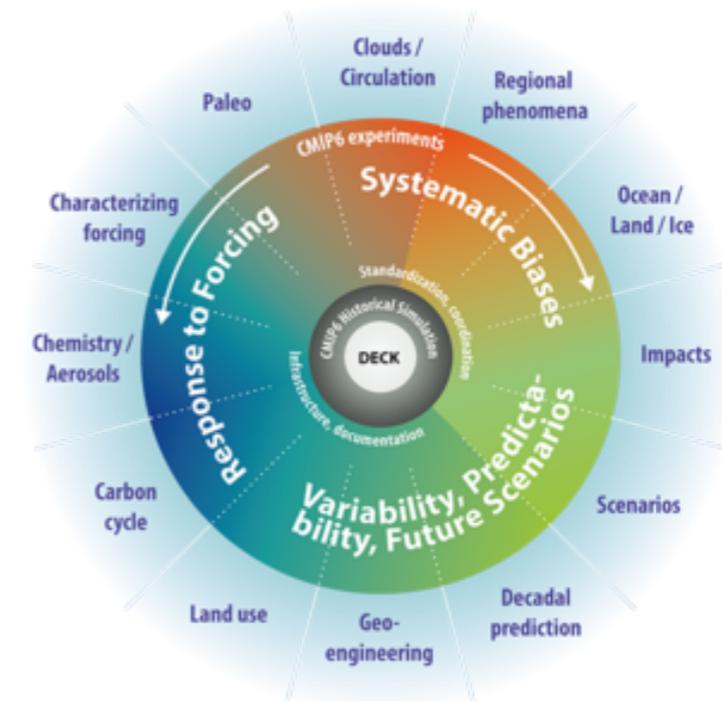
Sub-Task	Team Leads and Funding Agencies / Institution	Description
1. User Interface and Search and Dashboard Working Team	Luca Cinquini (JPL) and Guillaume Levavasseur (IPSL) and Alessandra Nuzzo (CMCC)	Improved ESGF search and data cart management and interface; ESGF search engine based on Solr5; discoverable search metadata; Statistics related to ESGF user metrics
2. Compute Working Team	Charles Doutriaux (LLNL) and Daniel Duffy (GSFC)	Developing the capability to enable data analytics within ESGF
3. Identity Entitlement Access Working Team	Philp Kershaw (CEDA) and Lukasz Lacinski (ANL)	ESGF X.509 certificate-based authentication and improved interface
4. Installation Working Team and Software Security Working Team	William Hill (LLNL), George Rumney (GSFC) and Sasha Ames (LLNL)	Installation of the components of the ESGF software stack; Security scans to identify vulnerabilities in the ESFF software;
5. International Climate Network Working Group, Replication/Versioning and Data Transfer Working Team	Eli Dart (ESnet), Lukasz Lacinski (ANL) and Stephan Kindermann (DKRZ)	Increase data transfer rates between the ESGF climate data centers; Replication tool for moving data from one ESGF center to another; ESGF data transfer and enhancement of the web-based download
6. Node Manager Working Team and Tracking / Feedback Notification Working Team	Sasha Ames (DOE), Tobias Weigel (DKRZ)	Management of ESGF nodes and node communications
7. Publication, Quality Control, Metadata, and Provenance Capture Working Team	Sasha Ames (LLNL), Katharina Berger (DKRZ) and Bibi Raju (DOE)	Capability to publish data sets for CMIP and other projects to ESGF; Integration of external information into the ESGF portal
8. User Support and Documentation Working Team	Matthew Harris (LLNL)	User frequently asked questions regarding ESGF and housed data; Document the use of the ESGF software stack
9. Containers Working Team	Luca Cinquini (JPL), Sebastien Gardoll (IPSL)	Design and implement a new ESGF architecture based on Docker containers and a selectable orchestration engine (Docker Swarm, Kubernetes, AWS ECS, etc...)

ESGF F2F 2017 PART 2

FUTURE CHALLENGES AND DEVELOPMENT PRIORITIES

UPCOMING SCIENTIFIC AND DATA CHALLENGES

- ▶ Support CMIP6 / IPCC AR6
 - ▶ 10X the CMIP5 data volumes (25-40 PB of uncompressed data)
 - ▶ Increased user community, number of ESGF nodes
 - ▶ Tens of related MIPs
- ▶ Integrate a larger number of observational data
 - ▶ Critical for validation and scoring of climate models
 - ▶ More Obs4MIPs data collections, data providers
- ▶ Scale performance of services to match the expected future data volumes
 - ▶ Searching, downloading, sub-setting, processing
 - ▶ Will require new architectures and more modularization
- ▶ Interoperate with other data infrastructures and agencies
 - ▶ NASA DAACs, NOAA, USGS, Copernicus, ACME...
- ▶ Provide reliable higher level data services
 - ▶ Efficient data exploration, analysis, visualization, learning
- ▶ Take advantage of commodity hardware and services, i.e. Cloud



DEVELOPMENT PRIORITIES IDENTIFIED BY 2016 F2F

- ▶ Easier to deploy and upgrade for system administrators
- ▶ Easier to use for end users
- ▶ Enable/expand server-side APIs and client toolkits to interact with the data programmatically
- ▶ Enable some basic server-side data processing
 - ▶ subsetting, regridding, averaging
- ▶ Proper citable identification of data via DOIs
- ▶ Improve overall system security
- ▶ Improve overall documentation



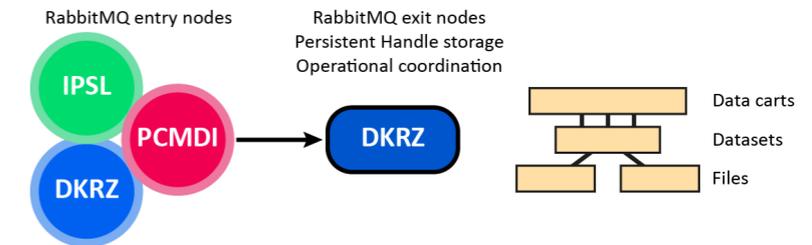
ESGF F2F 2017 PART 3

**SELECTED HIGHLIGHTS FROM
ONGOING TECHNICAL DEVELOPMENT**

CMIP6 DATA SERVICES

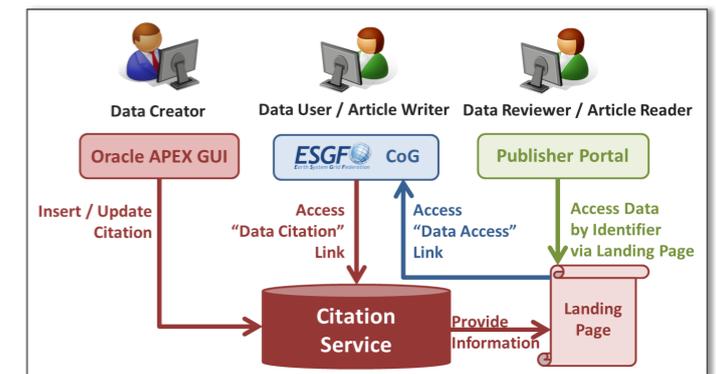
▶ Persistent Identifiers (PIDs)

- ▶ Provides persistent identifiers for all CMIP6 files and datasets
- ▶ Fully operational with installations at DKRZ, LLNL, IPSL



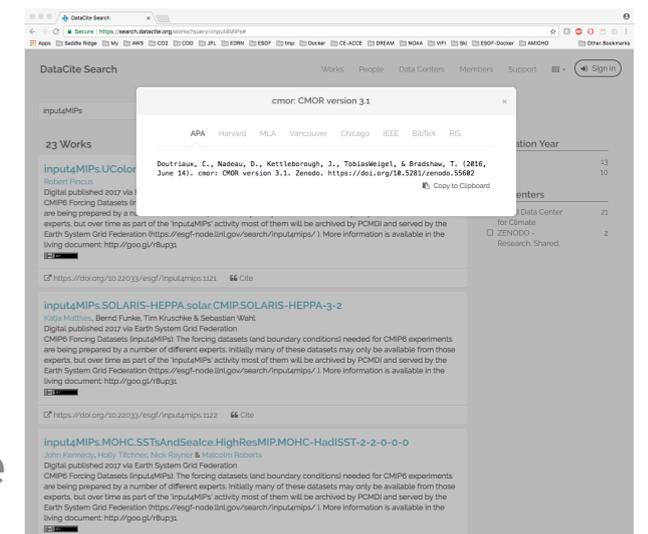
▶ Early Citation

- ▶ Provides citable identification for CMIP6 at model/MIP and experiment granularities before DOIs are assigned
- ▶ System in testing phase, operational release in 2018



▶ Digital Object Identifiers (DOIs)

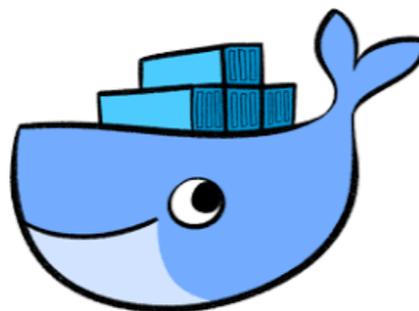
- ▶ Provides fully citable DOI at dataset level, after mandatory information has been populated
- ▶ CoG to display link to DOI landing pages at DKRZ, with information pushed to DataCite
- ▶ Citation GUI is fully functional, citation API under testing, to be released in early 2018



Acknowledgments: Stephan Kindermann (DKRZ), Martina Stockhouse (DKRZ)

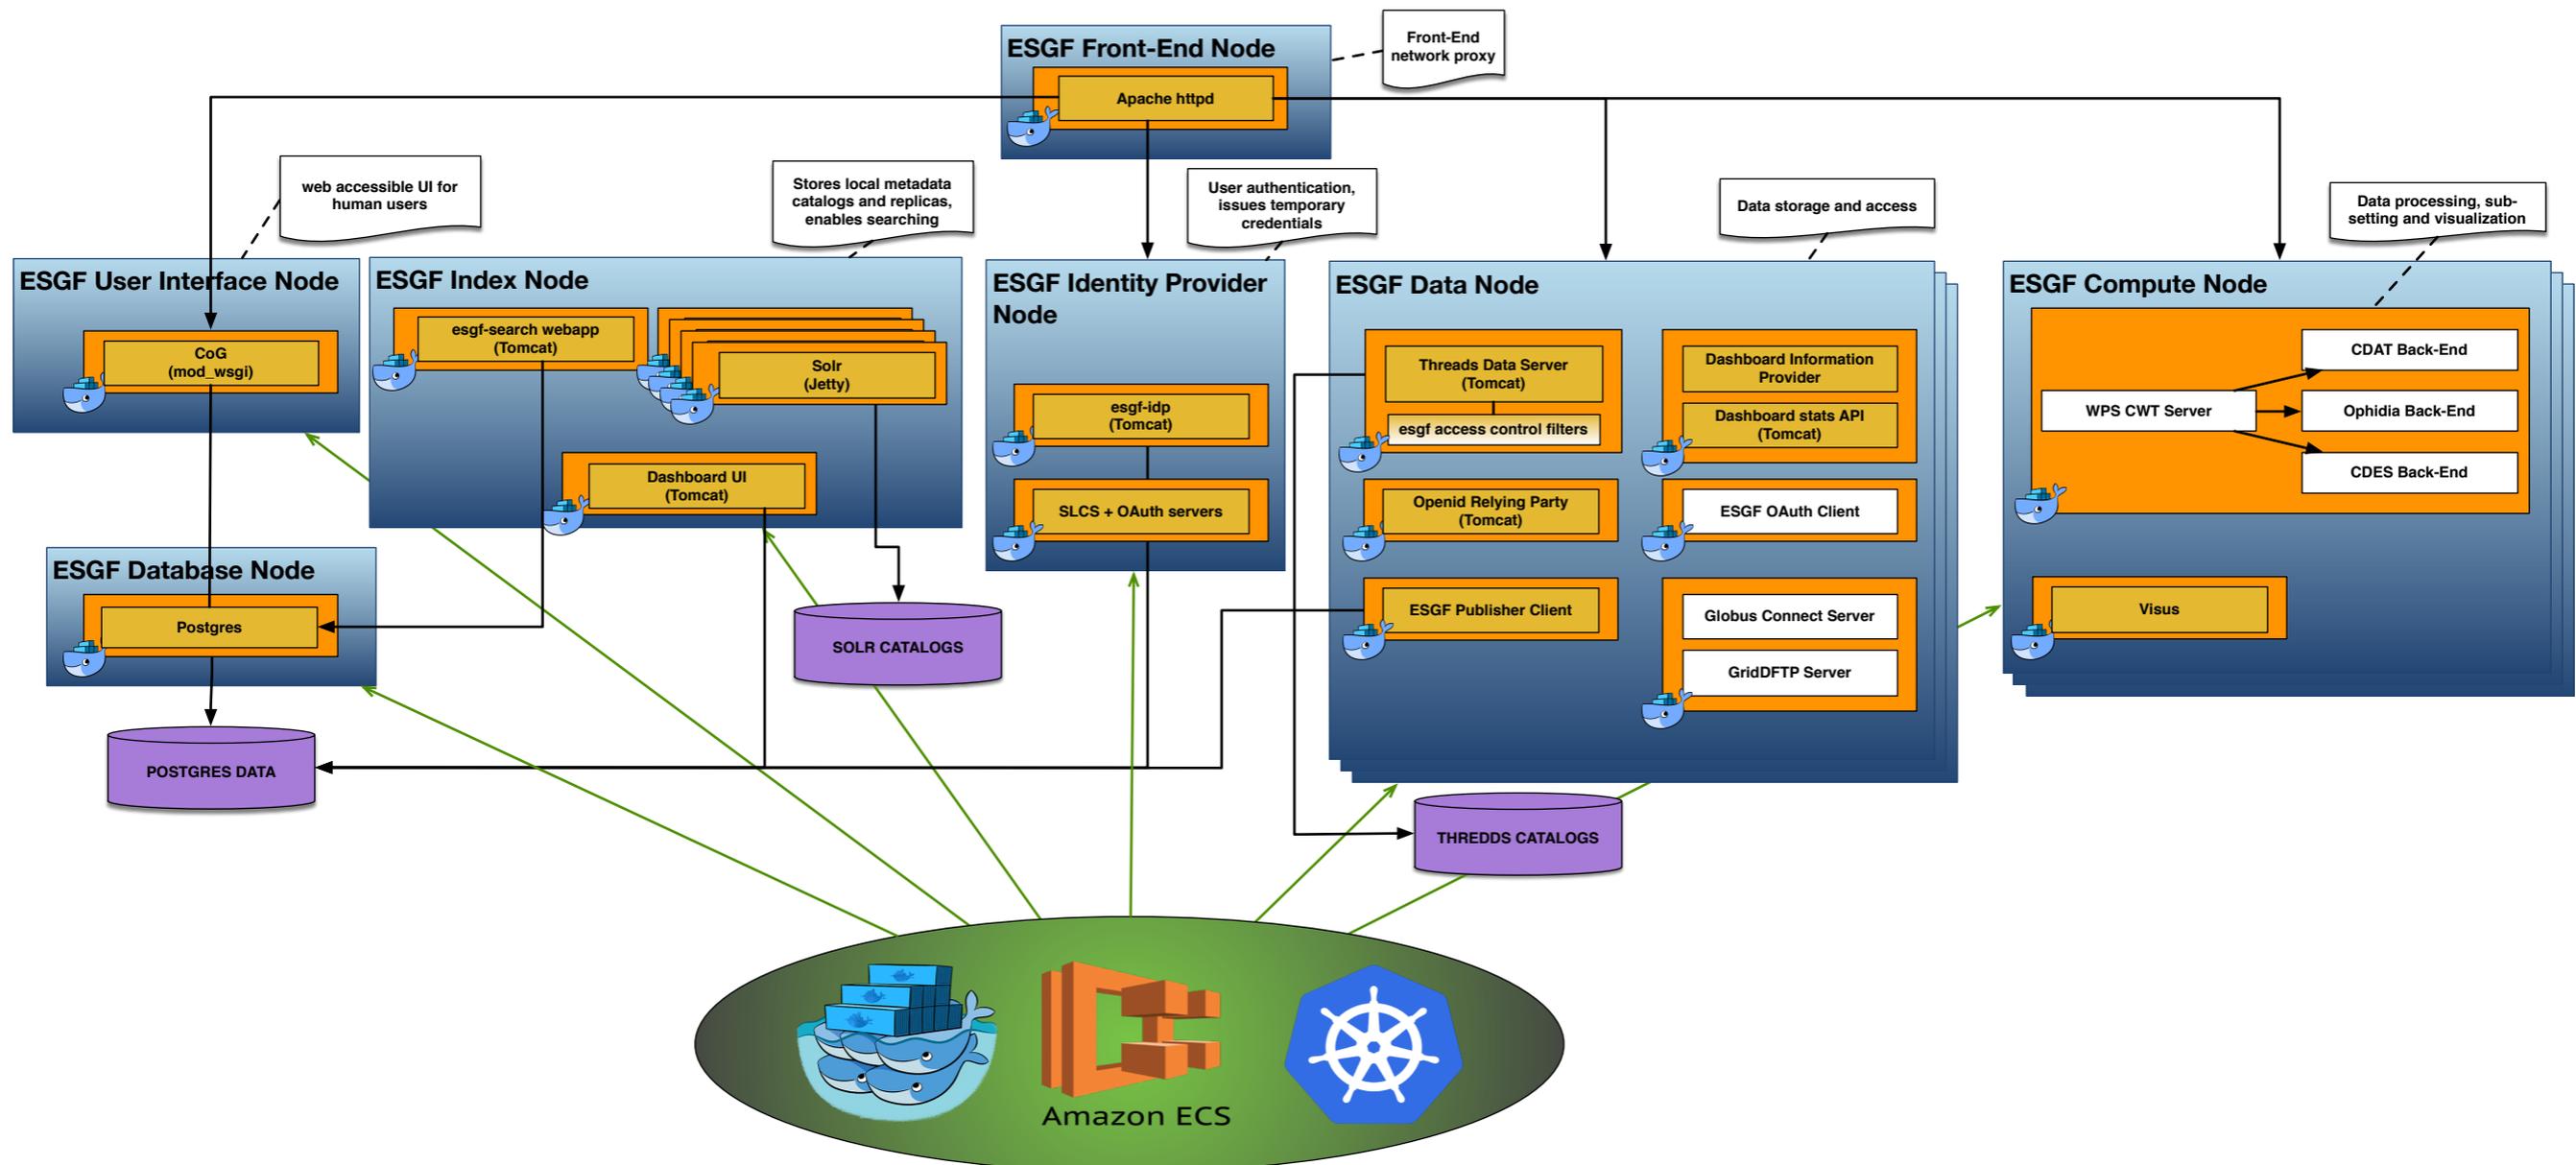
ESGF NODE INSTALLATION

- ▶ Traditionally, installation of an ESGF Node has been difficult to say the least. We are addressing this problem with 2 separate (synergistic) development efforts:
- ▶ Conversion of shell-based installer to Python based framework
 - ▶ More modular, easier to debug and extend, allow for independent installation and upgrade of single components
- ▶ Containerization of ESGF services
 - ▶ Evolution towards “micro-services” architecture
 - ▶ New architecture is based on Docker containers + different possible orchestration engines (Docker Swarm, AWS ECS, Kubernetes)
 - ▶ More modular and easier to install, more scalable, deployable anywhere



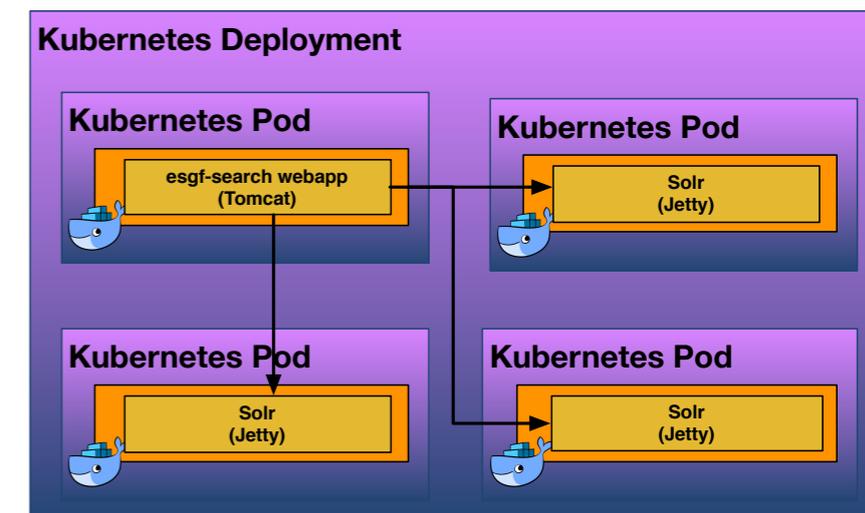
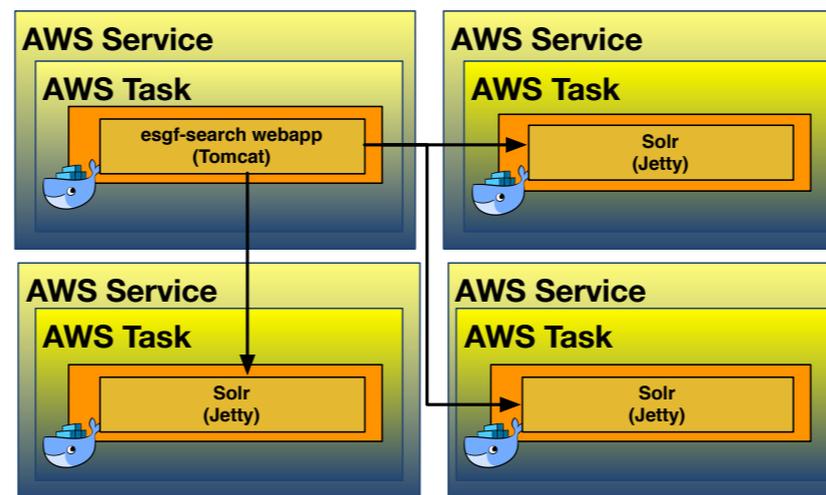
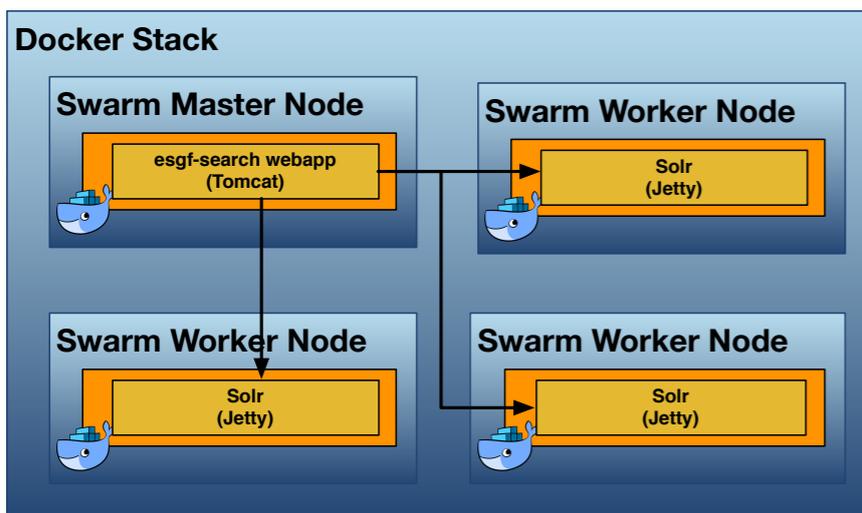
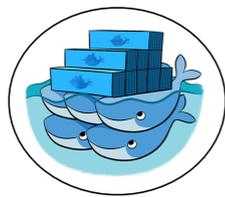
ESGF DOCKER ARCHITECTURE

- ▶ Each ESGF service is encapsulated within a Docker container
 - ▶ pre-built, can be tested, deployed, upgraded and scaled independently
- ▶ Containers are deployed onto a cluster of multiple Nodes with flexible metadata tags
 - ▶ multiple deployment configurations are possible
- ▶ Can scale the number of containers on each node, and the number of nodes of each kind
- ▶ Containers interact through orchestration engines: Docker Swarm, AWS ECS, Kubernetes..



ESGF ON THE CLOUD

- ▶ The ESGF Docker stack can be seamlessly deployed on the Cloud, exactly as it is deployed on an in-premise cluster
- ▶ ESGF+Cloud business model must still be worked out
 - ▶ There are costs associated with hardware, storage and networking (data in/out)
 - ▶ Particularly attractive for bursting computational tasks to volatile servers
- ▶ Experimented with Amazon Web Services + Docker Swarm or ECS, now adding Google Cloud with Kubernetes
- ▶ Goal: enable deployment of a fully operational ESGF Node on the Cloud by end of 2018



SERVER SIDE COMPUTING

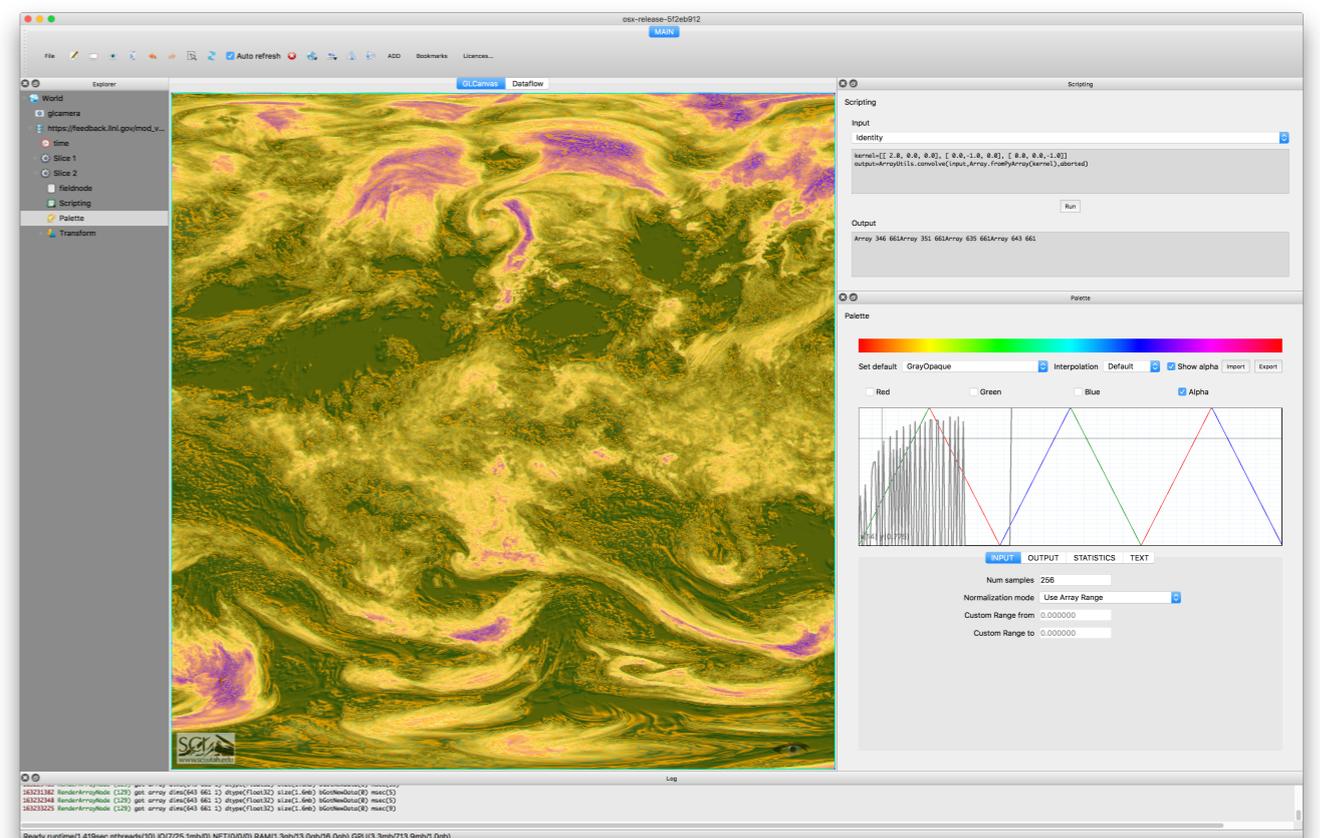
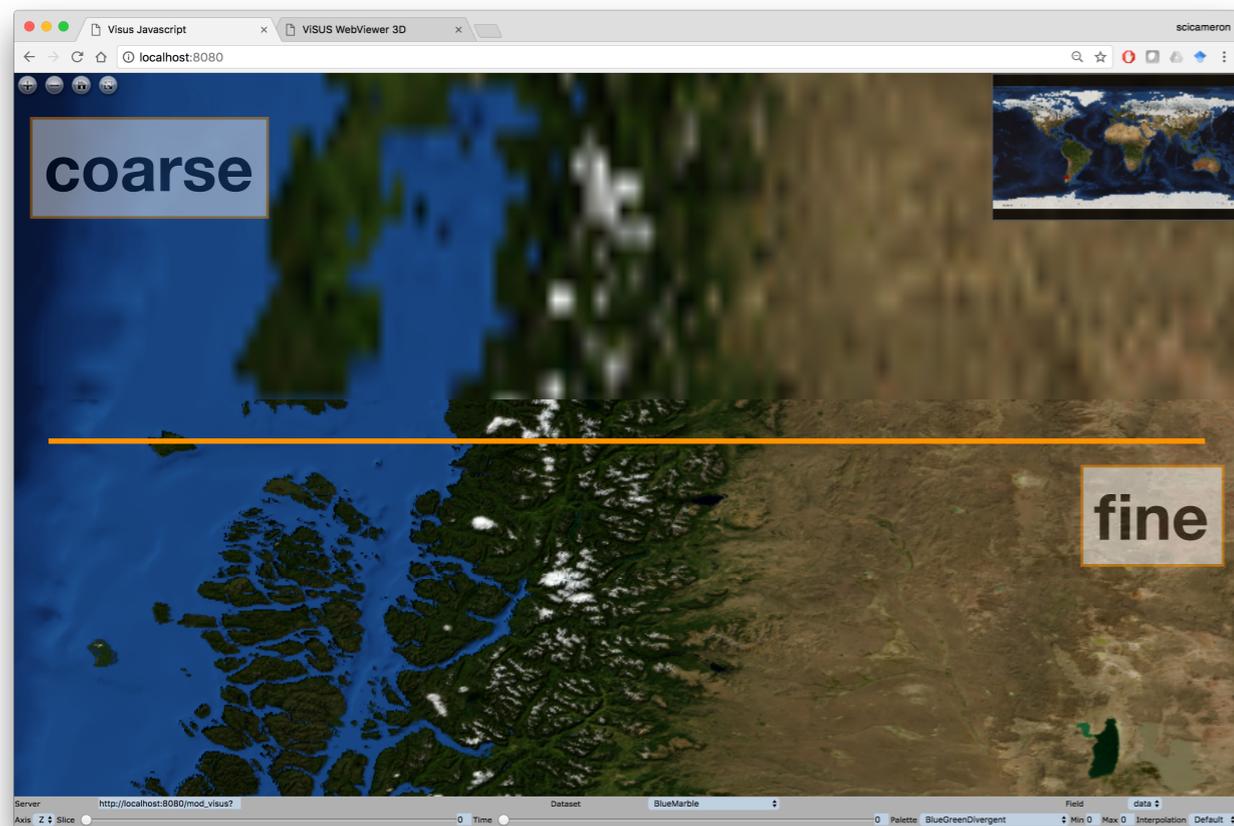
- ▶ Motivation: because of ever increasing data volumes, computations must be executed where data are stored, to avoid unnecessary network traffic
- ▶ Difficulty increased by distributed nature of ESGF archives
- ▶ CWT has been developing a client and server side API for distributed remote computing:
 - ▶ Based on OGC/WPS standard
 - ▶ Supports basic data operations (subsetting, regridding, averaging,...)
 - ▶ Multiple back-end implementations (CDAT, EDAS, Ophidia)
 - ▶ In beta, will start to be deployed at ESGF sites with next major release
 - ▶ Integrated with CoG UI

The screenshot displays the ESGF CoG (Climate Gateway) user interface. At the top, it shows the user is logged in as 'Admin' and is a 'LOCALHOST Node Administrator'. The main section is titled 'My Data Cart' and lists two datasets. The first dataset is 'project=CMIP5, model=NASA Goddard Institute for Space Studies, experiment=1 percent per year CO2, time_frequency=mon, modeling_real=atmos, ensemble=r11p2, version=20160425'. The second dataset is 'project=CMIP5, model=MIROC5, Atmosphere and Ocean Research Institute (The University of Tokyo), National Institute for Environmental Studies, and Japan Agency for Marine-Earth Science and Technology, experiment=10- or 30-year run initialized in year 2001, time_frequency=mon, modeling_real=atmos, ensemble=r41p1, version=20120710'. Below the datasets, there is a 'WPS Jobs' section with a 'Workflow' tab. The workflow interface includes a 'General' tab with fields for 'Dataset' (cmip5.output1.NCC.No), 'Variable' (clw), and 'Process' (CDAT.aggregate). There are buttons for 'Script' and 'Execute'. A map shows a blue rectangular selection box over the North Atlantic region. Below the map, there is a code editor with Python code for executing a WPS process and plotting the results. The code includes imports for 'cdms2' and 'vcs', and a plot of 'tas' (Near-Surface Air Temperature). The plot shows a color map of the North Atlantic region with a color scale from 285.5 to 307.767 K.

Acknowledgments: Jason Boutte (LLNL), Charles Doutriaux (LLNL)

STREAMING VISUALIZATION

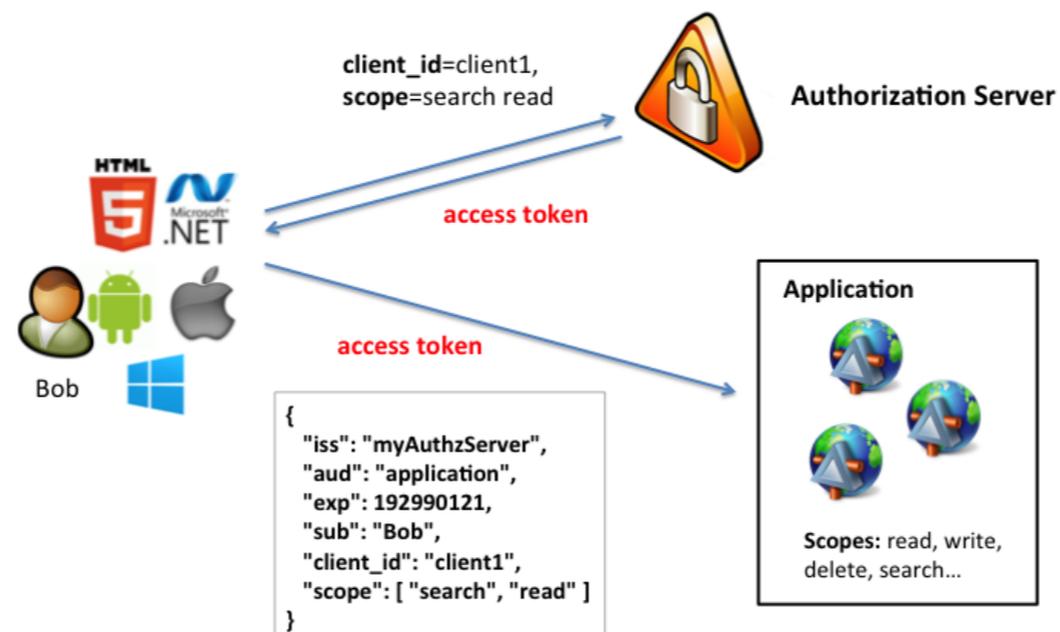
- ▶ Visus: new visualization engine developed by University of Utah as part of DREAM project
 - ▶ Streaming data at multiple image resolutions
 - ▶ Data converted on the fly from NetCDF to IDX
 - ▶ Supports scripting for server-side exploration and visualization of massive datasets



Acknowledgments: Cameron Christensen (University of Utah)

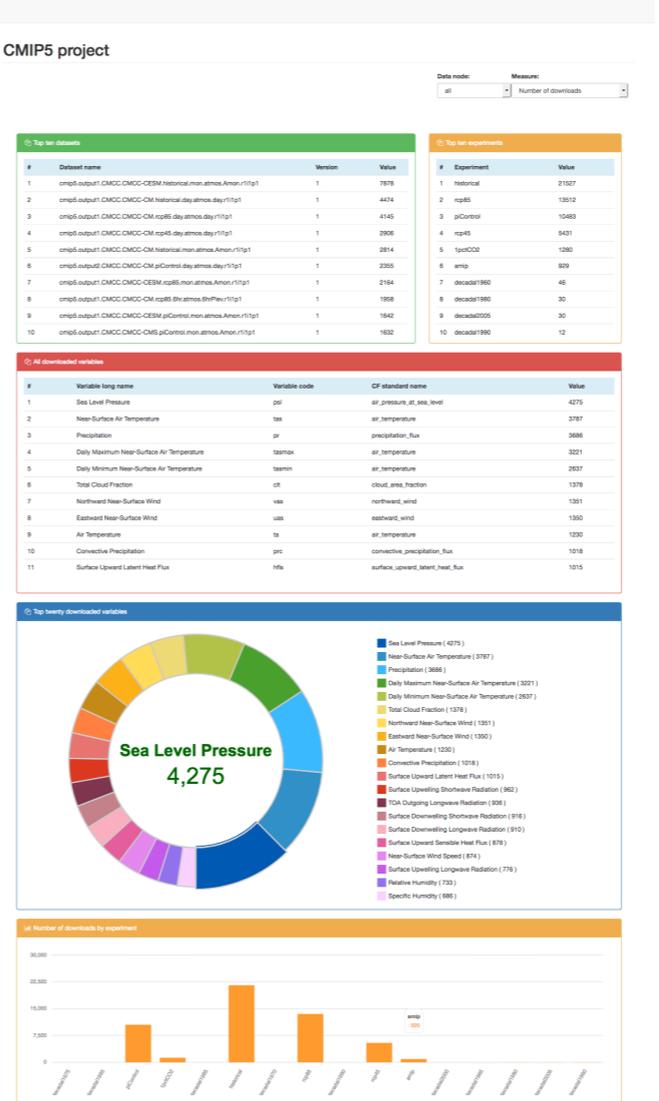
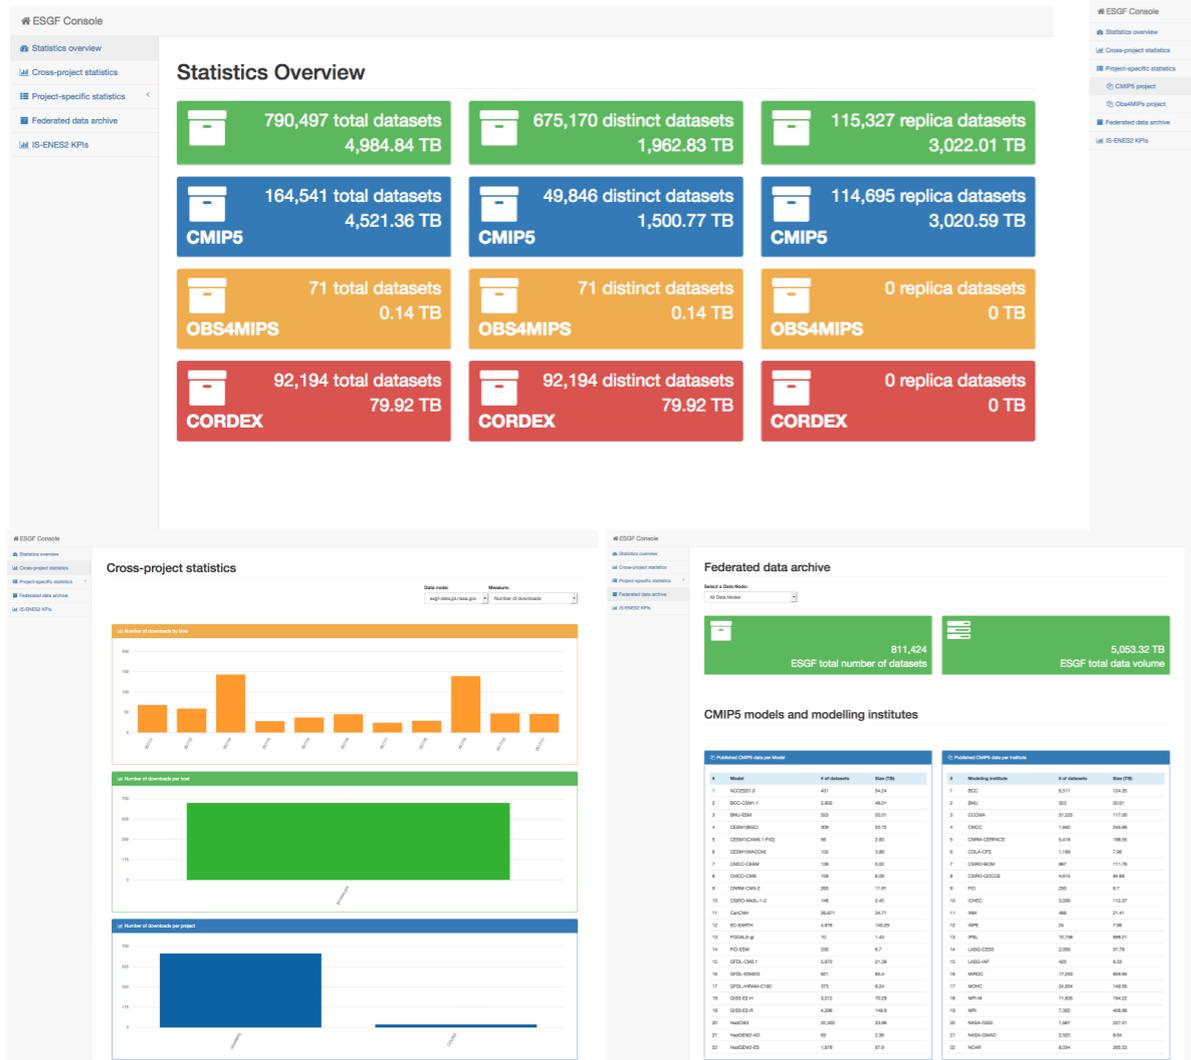
MIGRATION TO OAUTH2 BASED AUTHENTICATION

- ▶ Because ESGF is a federation of sites that are separately managed, the system needs both single-sign-on and distributed access control
- ▶ Current authentication protocol is based on Openid 2.0, which is now deprecated
- ▶ Several ESGF developers have been working on migrating the security infrastructure to OAuth2 and eventually OpenidConnect, which is the modern standard for authentication and authorization used by all industry enterprises
 - ▶ more secure, supported by modern community libraries
 - ▶ allows delegation of user credential for execution of server side processing
- ▶ ESGF OAuth2 services are in "alpha"
- ▶ ESGF plans to fully transition to the new system within 2018, without the users noticing...



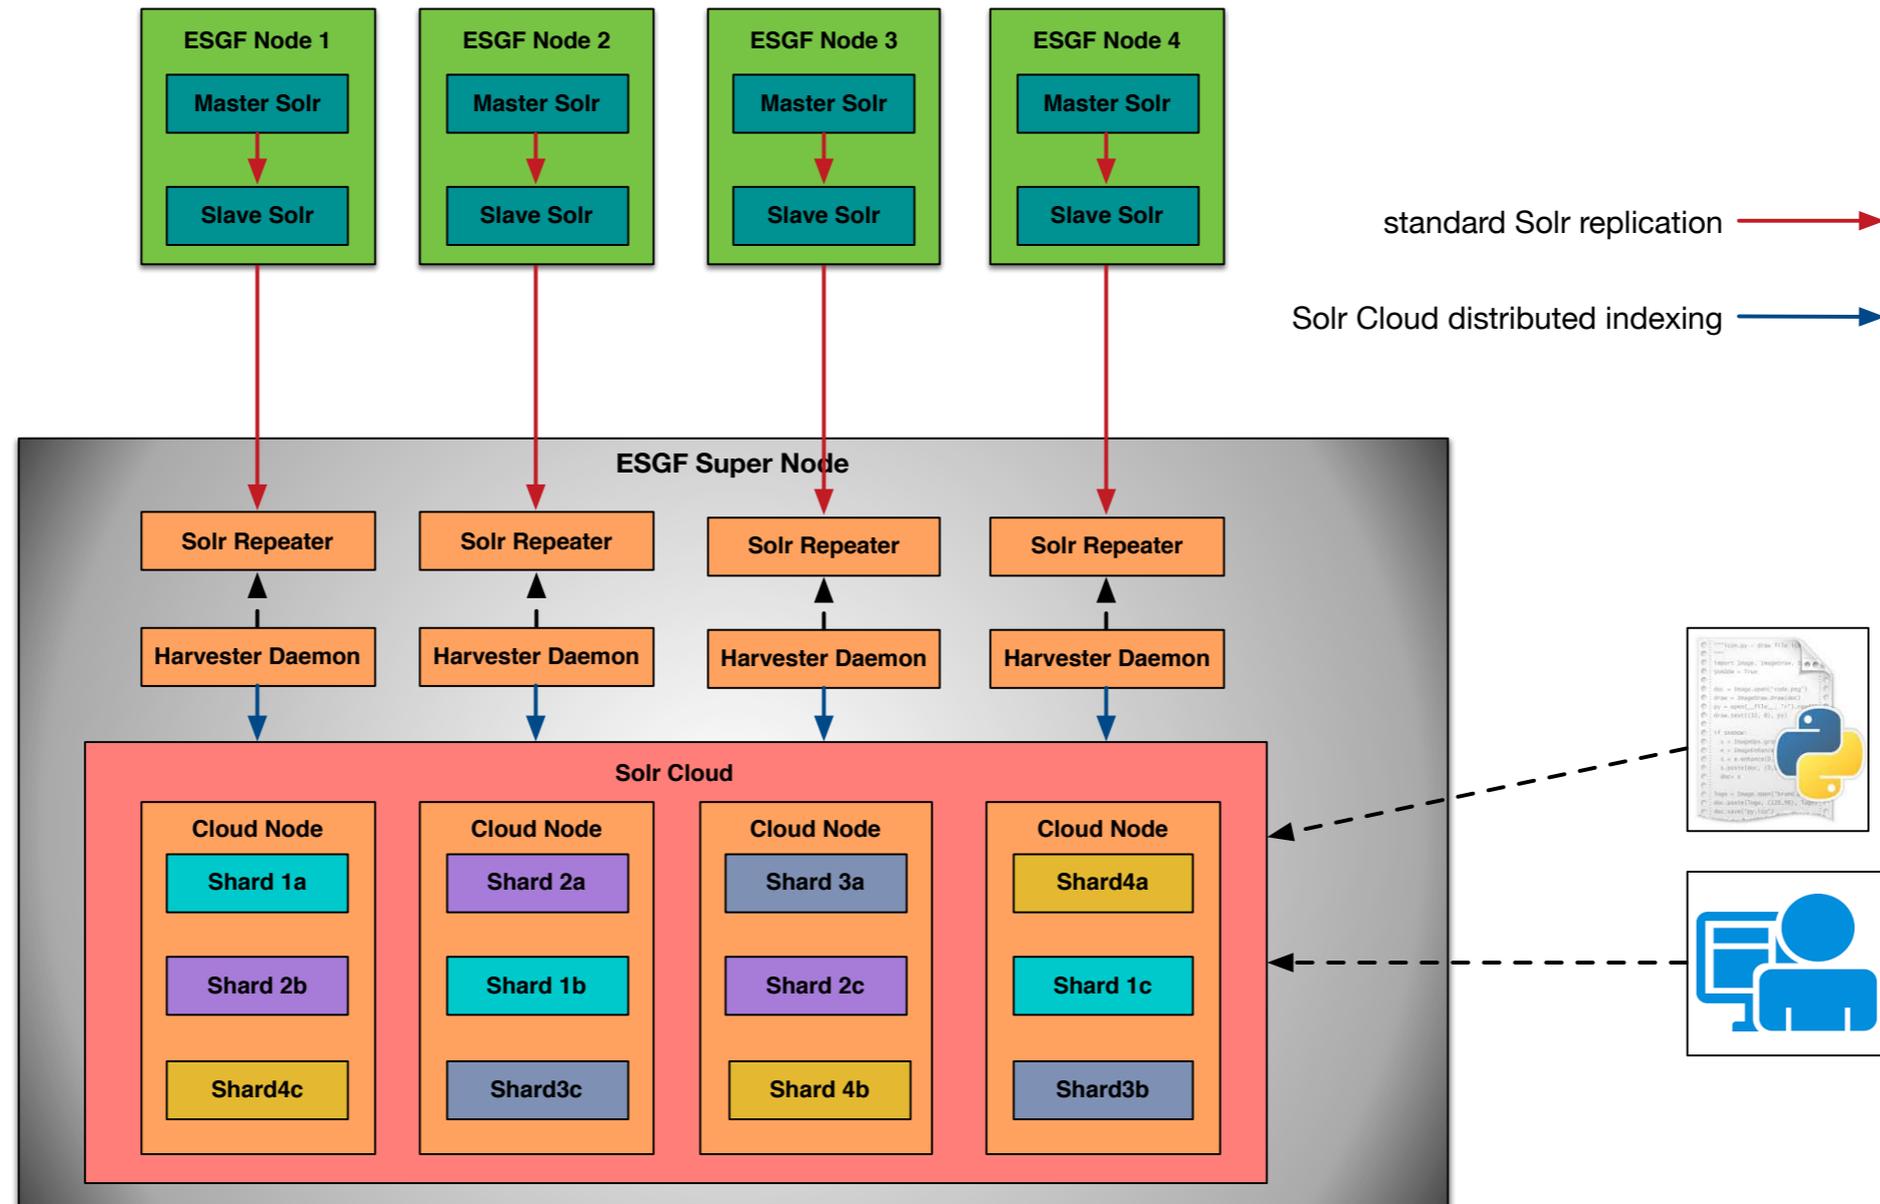
FEDERATED METRICS

- ▶ Dashboard team is finalizing the release of a fully functional dashboard package
 - ▶ collection of cross project and project specific metrics at each site
 - ▶ aggregation and presentation of federated metrics at selected collector sites
 - ▶ extensible model for defining new metrics and display views



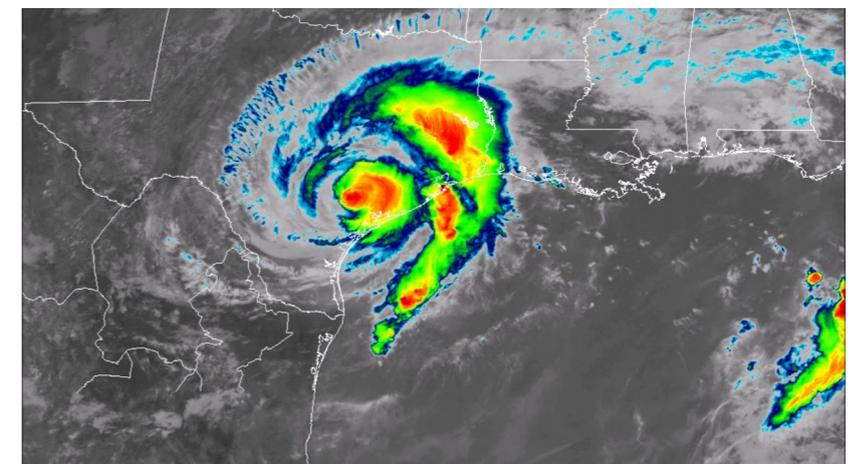
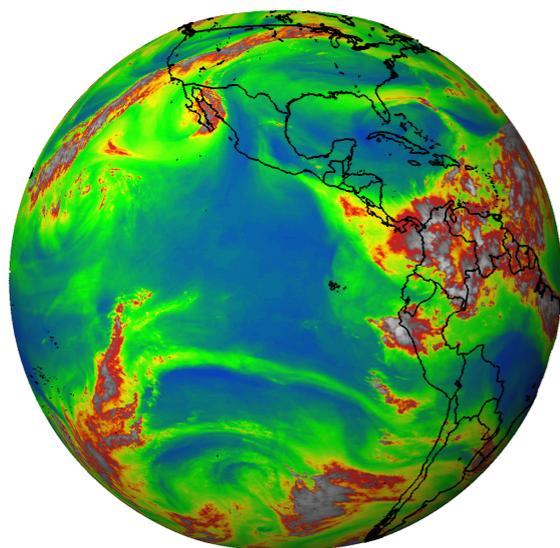
SCALABLE SEARCH

- ▶ Use Solr-Cloud to host a "super-search-node", possibly on the Cloud, for fastest search performance
- ▶ Prototype architecture:
 - ▶ Solr repeaters to replicate information from individual sites via standard replication
 - ▶ Harvester daemons to transfer metadata to Solr-Cloud nodes and shards
- ▶ Advantages: automatic distributed indexing and searching, high availability, automatic failover, enhanced performance, ...



OTHER DEVELOPMENT

- ▶ Diagnostics: how do we thoroughly validate the model data that are generated, or the observational data that are collected, so that scientists can use them with confidence ?
- ▶ Provenance: how do we completely capture the history (inputs and algorithms) of a specific science result, so that it is reproducible and justifiable ?
- ▶ Machine Learning: how do we automatically detect features and patterns from the vast distributed archives that we manage ?
- ▶ ...and much more...



ESGF 2017 F2F AGENDA SUMMARY

- ▶ Day 1 (Tuesday 12/5):
 - ▶ Science Drivers
 - ▶ ESGF Awards Ceremony
 - ▶ Posters and Demos
- ▶ Day 2 (Wednesday 12/6):
 - ▶ Working Groups Presentations and Plenary Discussions
 - ▶ Invited Google Cloud Presentation
- ▶ Day 3 (Thursday 12/7):
 - ▶ Working Groups Presentations and Plenary Discussions
 - ▶ CDNOT ("CMIP6 Data Nodes Operations Team") Readiness Discussion
- ▶ Day 4 (Friday 12/8):
 - ▶ Executive Committee & Working Groups Breakout Meetings
 - ▶ Executive Committee & Working Groups Conference Findings and Final Remarks
 - ▶ Optional Code Sprint

KEY QUESTIONS FOR 2017 F2F

- ▶ Is ESGF ready to manage and serve CMIP6 data to the climate community ?
- ▶ Can ESGF last longer than a few years, proving that it can scale to the future larger data volumes of climate models and observations ?
- ▶ How can ESGF be made easier to use for node administrators and end-users ?
- ▶ What should be the ESGF priorities for the next 3-5 years ?

