



Partnerships for development of next-generation software for distributed access and analysis of simulated, observed, and reanalysis data from the climate and weather communities.

**Joint DOE, NASA, NOAA, NSF, IS-ENES, and ANU/NCI
Earth System Grid Federation (ESGF)
2017 Face-to-Face Conference Abstracts**

**Day 1: Tuesday, 5 December 2017
ESGF Steering Committee and Executive Committee**

Title and Presenter	Abstract
<p>Department of Energy Office of Biological and Environmental Research Data Management</p> <p>Opening Comments</p> <p>Justin Hnilo (DOE/BER) <i>Justin.Hnilo@science.doe.gov</i></p>	
<p>The State of the Earth System Grid Federation</p> <p>Dean N. Williams (DOE/LLNL/AIMS) <i>williams13@llnl.gov</i></p>	<p>The Earth System Grid Federation (ESGF) is a multi-institutional, international software infrastructure and development collaboration led by scientists and software engineers worldwide. ESGF’s mission is to facilitate scientific research and discovery on a global scale. The ESGF architecture federates a geographically distributed network of climate modeling and data centers that are independently administered yet united by common protocols and application programming interfaces. The cornerstone of its interoperability is peer-to-peer messaging, which continuously exchanges information among all nodes through a shared, secure architecture for search and discovery. The ESGF integrates popular open-source application engines with custom components for data publishing, searching, user interface, security, metrics, and messaging to provide petabytes of geophysical data to roughly 25,000 users from over 1,400 sites on six continents. It contains output from the Coupled Model Intercomparison Project (CMIP), used by authors of the Intergovernmental Panel on Climate Change Third, Fifth, and Sixth Assessment Reports, and output from DOE’s Accelerated Climate Modeling for Energy (ACME) project and the European Union’s Copernicus Programme. Over the next three years, we propose to (1) sustain and enhance a resilient data infrastructure with friendlier tools for the expanding global scientific community and (2) prototype new tools that fill important capability gaps in scientific data archiving, access, and analysis. These goals will support a data-sharing ecosystem and, ultimately, provide predictive understanding of couplings and feedbacks among natural-system and anthropogenic processes across a wide range of geophysical spatial scales.</p>

Day 1: Tuesday, 5 December 2017

ESGF Science Drivers: Project Requirements and Feedback

Title and Presenter	Abstract
<p>Coupled Model Intercomparison Project, phase 6 (CMIP6) and the Working Group on Coupled Modeling Infrastructure Panel (WIP)</p> <p>Karl Taylor (DOE/LLNL/PCMDI) taylor13@llnl.gov</p>	
<p>Observations for Model Intercomparison Project (Obs4MIPs)</p> <p>Peter Gleckler (DOE/LLNL/PCMDI) gleckler1@llnl.gov</p> <p>Robert Ferraro (NASA/JPL) robert.d.ferraro@jpl.nasa.gov</p>	
<p>Copernicus Programme</p> <p>Sébastien Denvil (ENES/IPSL) sebastien.denvil@ipsl.jussieu.fr</p>	
<p>Collaborative REAnalysis Technical Environment</p> <p>Jerry Potter (NASA/GSFC) gerald.potter@nasa.gov</p>	
<p>Accelerated Climate Model for Energy (ACME) Workflow</p> <p>Dean N. Williams (DOE/LLNL/AIMS) williams13@llnl.gov</p>	<p>The advanced model development, testing, and execution infrastructure has been designed to strongly accelerate the model development and testing cycle for the new Department of Energy (DOE) Accelerated Climate Model for Energy (ACME) model by automating labor-intensive tasks, providing intelligent support for complex tasks, and reducing duplication of effort through collaborative Workflow Group support. The Workflow Group had two important assignments: (1) advance model development by developing, testing, and executing an end-to-end infrastructure that automates labor-intensive tasks; and (2) provide intelligent support for complex tasks in model development through scientific model component (i.e., atmosphere, land, ocean, and sea ice) collaboration.</p> <p>To achieve our primary objectives, the team was split into several epic subtasks: (1) ACME Workbench and Process Flow; (2) Data Management; (3) Analysis and Visualization; (4) Diagnostics; (5) Provenance Capture; and (6) Hardware Infrastructure. These open-source projects have grown in scope as requirements have shifted or completely changed over the course of the project. The tools and experience resulting from their development provides the foundation on which the end-to-end model test bed infrastructure will be based. As the global view of the ACME project expands across the component model space, the usefulness and urgency of the workflow software becomes more apparent. The end goal of every quarter for the Workflow Group was to advance a step closer to reducing the level of effort to successfully run the ACME model, archive output, generate diagnostics, and share the results of both the model output and diagnostics results</p>

Title and Presenter	Abstract
	with ACME colleagues.

Day 1: Tuesday, 5 December 2017
Poster and Live Demonstration Session

Title and Presenter	Abstract
<p style="text-align: center;">Poster 1</p> <p>Name and Affiliation email@address.agency</p>	
<p style="text-align: center;">Poster 2</p> <p>Name and Affiliation email@address.agency</p>	
<p style="text-align: center;">Poster 3</p> <p>Name and Affiliation email@address.agency</p>	
<p style="text-align: center;">Poster 4</p> <p>Name and Affiliation email@address.agency</p>	
<p style="text-align: center;">Poster 5</p> <p>Name and Affiliation email@address.agency</p>	
<p style="text-align: center;">Poster 6</p> <p>Name and Affiliation email@address.agency</p>	
<p style="text-align: center;">Demo 1</p> <p>Name and Affiliation email@address.agency</p>	
<p style="text-align: center;">Demo 2</p> <p>Name and Affiliation email@address.agency</p>	
<p style="text-align: center;">Demo 3</p>	

Title and Presenter	Abstract
Name and Affiliation email@address.agency	
Demo 4 Name and Affiliation email@address.agency	
Demo 5 Name and Affiliation email@address.agency	
Demo 6 Name and Affiliation email@address.agency	

Day 2: Wednesday, 6 December 2017

ESGF Focus Areas

Title and Presenter	Abstract
International Climate Network Working Group, Replication / Versioning and Data Transfer Working Team Plenary Eli Dart (DOE/ESnet) dart@es.net Lukasz Lacinski (DOE/ANL) lukasz@uchicago.edu Stephan Kindermann (ENES/DKRZ) kindermann@dkrz.de	
Compute and Data Analytics Working Team Plenary <i>Charles Doutriaux (DOE/LLNL/AIMS)</i> doutriaux1@llnl.gov Daniel Duffy (NASA/GSFC) daniel.q.duffy@nasa.gov	
Identity Entitlement Access Working Team Plenary	

Title and Presenter	Abstract
<p>Phil Kershaw (ENES/CEDA) <i>philip.kershaw@stfc.ac.uk</i></p> <p>Lukasz Lacinski (DOE/ANL) <i>lukasz@uchicago.edu</i></p>	
<p>Status update and future planning for the ESGF UI, Search, and Dashboard Working Group</p> <p>Luca Cinquini (NASA/JPL) <i>Luca.Cinquini@jpl.nasa.gov</i></p> <p>LEVAVASSEUR Guillaume (ENES/IPSL) <i>glipsl@ipsl.jussieu.fr</i></p> <p>Alessandra Nuzzo (ENES/CMCC) <i>alessandra.nuzzo@cmcc.it</i></p>	<p>This presentation will provide a progress report and future roadmap for the recently unified working group that includes the CoG UI, the search back-end services, and the Dashboard and metrics functionality.</p> <p>CoG development has been focused on integrating new features in support of critical community projects such as the upcoming CMIP6 and the ongoing Obs4MIPs. If future funding is provided, we plan to completely re-factor the CoG software to enhance its modularity, functionality, security, and extensibility. Also, because the front-end is more and more requested by non-scientific users from different backgrounds, future efforts must lead to a friendlier interface with intuitive layouts and helpful tutorials.</p> <p>The back-end search services have been mostly stable, with some development effort again focused on supporting CMIP6 features, as well as addressing newly discovered security vulnerabilities. Future work must address necessary Solr upgrades and perhaps moving to Solr Cloud.</p> <p>The Dashboard application has been completely re-written since the ESGF shut-down, and is now deployed as an information provider as part of each Data Node, including a RESTful API. Federation-level metrics are provided by aggregator applications that will be deployed at selected Tier 1 sites.</p>
<p>Installation and Software Security Working Team Plenary</p> <p>William Hill (DOE/LLNL/AIMS) <i>hill119@llnl.gov</i></p> <p>Prashanth Dwarakanath (ENES/Liu) <i>pchengi@nsc.liu.se</i></p> <p>Luca Cinquini (NASA/JPL) <i>Luca.Cinquini@jpl.nasa.gov</i></p> <p>George Rumney (NASA/GSFC) <i>george.rumney@nasa.gov</i></p>	<p>William Hill and Prashanth Dwarakanath – Software Installation</p> <p>Abstract:</p> <hr/> <p>Luca Cinquini – Software Container (i.e., Docker)</p> <p>Abstract:</p> <p>This presentation will report on the current state of the effort to design and implement a "next-generation" ESGF architecture based on Docker containers. Such a model presents great advantages with respect to the current "monolithic" architecture supported by the shell-based installer, such as easier to install and upgrade, scalable onto multiple hosts, deployable both on internal clusters and commercial Cloud. This work has been so far supported by the DREAM project, and is now joining forces with the new European Copernicus project.</p> <hr/> <p>George Rumney – Software Security</p> <p>Abstract:</p>

Day 3: Thursday, 7 December 2017

Coordinated Efforts with Community Software Projects

Title and Presenter	Abstract
<p>Publication, Quality Control,</p>	<p>Sasha Ames – Publication Progress</p>

Title and Presenter	Abstract
<p>Metadata, and Provenance Capture Working Team Plenary</p> <p>Sasha Ames (DOE/LLNL/AIMS) <i>ames4@llnl.gov</i></p> <p>Heinz-Dieter Hollweg (ENES/DKRZ) <i>hollweg@dkrz.de</i></p>	<p>Abstract:</p> <hr/> <p>Heinz-Dieter Hollweg – Quality Control Progress</p> <p>Abstract:</p>
<p>Machine Learning Plenary</p> <p>Sookyung Kim (DOE/LLNL/AIMS) <i>kim79@llnl.gov</i></p>	<p>Sookyung Kim – Community Machine Learning</p> <p>Abstract:</p> <hr/> <p>Copernicus Presenter – Title</p> <p>Abstract:</p>
<p>Diagnostics Plenary</p> <p>Zeshawn Shaheen (DOE/LLNL/AIMS) <i>shaheen2@llnl.gov</i></p>	<p>Zeshawn Shaheen – Community Diagnostics Package</p> <p>Abstract:</p> <hr/> <p>Tom Landry – Canada Diagnostics</p> <p>Abstract:</p> <hr/> <p>Copernicus Presenter – Title</p> <p>Abstract:</p>
<p>Cloud Plenary</p> <p>Luca Cinquini (NASA/JPL) <i>Luca.Cinquini@jpl.nasa.gov</i></p> <p>V. Balaji (NOAA/GFDL) <i>balaji@princeton.edu</i></p> <p>Tom Landry (CRIM) <i>tom.landry@crim.ca</i></p>	<p>Luca Cinquini – Amazon Cloud and Collaboration</p> <p>Abstract:</p> <hr/> <p>V. Balaji – Google Cloud and Collaboration</p> <p>Abstract:</p> <hr/> <p>Tom Landry – Microsoft Cloud and Collaboration</p> <p>Abstract:</p>
<p>Node Manager and Tracking / Feedback Notification Plenary</p> <p>Sasha Ames (DOE/LLNL/AIMS) <i>ames4@llnl.gov</i></p> <p>Tobias Weigel (ENES/DKRZ) <i>weigel@dkrz.de</i></p>	<p>Sasha Ames – Node Manager</p> <p>Abstract:</p> <hr/> <p>Tobias Weigel – PID Services and Tracking/Feedback</p> <p>Abstract:</p>
<p>User Support and Documentation Plenary</p>	

Title and Presenter	Abstract
<p>Matthew Harris (DOE/LLNL/AIMS) <i>harris112@llnl.gov</i></p>	