NASA Science Mission Directorate
Research Opportunities in Space and Earth Sciences – 2015
NNH15ZDA001N-ACCESS
A.36 Advancing Collaborative Connections for Earth System Science

The Advancing Collaborative Connections for Earth System Science (ACCESS) Program solicits projects to leverage existing technologies that enhance and improve data and information systems in support of NASA's Earth science research goals. The 2015 ACCESS announcement specifically sought technologies that improve data discovery, accessibility and usability for users of NASA's Earth science data systems, emphasizing interoperability and harmonized user experience.

A total of 52 proposals were received for this announcement. All proposals were peer evaluated using an expert panel review. The Earth Science Division of NASA's Science Mission Directorate selected 10 proposals for two-year awards pending satisfactory budget and work plan negotiations. The total first-year funding for these proposals is approximately 3.5 million dollars. These projects will help to further improve NASA's Earth Science Division's heterogeneous and distributed data and information systems.

Yehuda Bock/University of California, San Diego
Deploying Technology for Distributed Use of Global Navigation Satellite System Products in Earthquake and Tsunami Warning
15-ACCESS15-0045

Rapid modeling of the earthquake source is critical to earthquake and tsunami early warning systems, in particular for mitigating hazards to near-source coastal communities. Operational earthquake early warning and rapid response in the U.S. is the role of the United States Geological Survey, and for tsunamis, the National Oceanic and Atmospheric Administration through its Tsunami Warning Centers (TWC). Traditionally, monitoring of earthquakes and tsunamis has been based on seismic networks for rapidly estimating earthquake magnitude and slip, and tide gauges and deep-ocean buoys for measurement of tsunami waves offshore. Since the under-estimation of the size of the 2011 Mw 9.0 Tohoku-Oki, Japan, earthquake and the resulting under-prediction of tsunami inundation in the presence of the most extensively instrumented warning system in the world, it has become clear that additional observations are necessary to mitigate a future repeat of the massive destruction as a consequence of large subduction zone earthquakes. Replay of the recorded data from Tohoku and other large subduction zone events has definitively demonstrated that now-time analysis of high-rate GPS, and more generally GNSS, observations for station displacements provides rapid magnitude estimates and direct fault displacement measurements that remove any ambiguity about the maximum size of the earthquake. It provides direct inputs to tsunami generation models, especially critical for the most vulnerable populations and infrastructure nearest the earthquake source. The integration of GNSS with other data types, i.e., seismogeodesy (the combination of GNSS and strong-motion accelerometers), near-shore GNSS buoys and ocean-bottom pressure sensors has proven to further increase warning accuracy.

This proposal leverages more than a decade of research funded by NASA and NSF at the Scripps Institution of Oceanography's Orbit and Permanent Array Center (SOPAC) and other research groups in developing prototype geodetic-enhanced earthquake and tsunami early warning systems. Based on our development of algorithms for each of the required functions, we will implement plug-in modules for real-time access to and management of GNSS data, estimation of displacement and velocity waveforms using precise point positioning (PPP) technology, rapid determination of earthquake source properties, sea floor uplift parameters, and tsunami extent and inundation using seamless interfaces to a user's other in-house Earth observations and models. Specifically, we will collaborate with NOAA TWCs to integrate now-time GNSS-based observations and models into their operational systems. Collaboration with NOAA on tsunami warning in the Indo-Pacific Rim is a key NASA objective and falls within the Earth Surface and Interior science focus area "How can our knowledge of Earth surface change be used to predict and mitigate natural hazards?"

The primary objective of our proposal is seamlessly enabling GNSS into NOAA TWC operations with the following Java and C based plug-ins:

(1) GNSS real-time data streaming in parallel with other data types.
(2) GNSS PPP with ambiguity resolution with the option to optimally integrate collocated accelerometer data.
(3) Near-source real-time earthquake detection and rapid magnitude estimation.
(4) Earthquake fault parameter estimation as input to sea floor uplift models.
(5) Tsunami prediction methodology based on heterogeneous data.

Considering the maturity of SIO's system, the project will begin at TRL 7 and end at TRL 8 after two years. The links with the respective agencies have already been established through the Real-Time Earthquake Analysis for Disaster mItigation network (READI) working group and requests for these capabilities have originated within NOAA. Therefore, we expect that the operations concept will be supported beyond the two-year ACCESS funding and become operational at NOAA (TRL 9).

Adrian Borsa/University of California, San Diego
Advanced Discovery, Processing, and Visualization Services for ICESat and ICESat-2 Altimeter Data
15-ACCESS15-0003

Current approaches for the discovery and distribution of airborne and satellite remote sensing data are targeted primarily toward raster and/or point cloud datasets. These data are typically distributed on the web via a map-based GUI, with background optical imagery or shaded topography providing the spatial context for the data. Remote sensing measurements from profiling instruments such as laser/radar altimeters and atmospheric sounders present a different challenge. These observations are densely sampled along ground tracks that can be separated by kilometers or more and are largely independent of each other spatially. Profiling data are also information-rich, and often include vector (rather than scalar) measurements and many auxiliary parameters.

There are currently no open data access systems designed to address two critical user requirements for profiling data: 1) the ability to visualize vector data and auxiliary parameters along one or more individual tracks on a web-based map interface, and 2) simultaneous visualization of repeated tracks at the same location over time. NASA's ICESat and upcoming ICESat-2 missions use profiling laser altimeters to measure changes in the topography of Earth's ice sheets, vegetation canopy structure, and clouds and aerosols. The unique data from these missions require a new paradigm for data access, to serve the needs of a diverse scientific community and to increase the accessibility and utility of these data for new users.

We propose to build a cyberinfrastructure platform for ICESat and ICESat-2 data discovery, access, and visualization. This platform will leverage and build upon the service-oriented architecture behind the OpenTopography system (TRL8+) that already provides integrated access and processing capabilities for high resolution topography data. The proposed work will also leverage development performed as part of the NASA Lidar Access System project, a collaboration between UNAVCO, OpenTopography, the Goddard Space Flight Center, and the National Snow and Ice Data Center (NSIDC) that prototyped an ICESat data discovery interface and improved GLAS data products.

This cyberinfrastructure platform will include:

1. An efficient, user-friendly map-based web interface to access ICESat (single beam) and ICESat-2 (multiple-beam) data, with the ability to subset by time and by track as well as to filter data on multiple engineering and data quality attributes.

2. Visualization of the waveform/energy (ICESat-1) or photon (ICESat-2) distribution through the entire atmosphere and (separately) within the ground range window. In both cases, users would be able to visualize individual laser shots or horizontal stacks of multiple shots along profile.

3. An efficient and highly scalable data management platform that includes intelligent tiered storage and distributed file systems for redundancy and fault tolerance.

4. A standards-based API for enabling programmatic access to the data and processing capabilities for expert users and interoperability with other systems.

Profiling altimeter data are relevant to 5 of the 6 interdisciplinary science focus areas for NASA, with important applications in atmospheric science, terrestrial hydrology, sea level change, vegetation/biomass monitoring, and the cryospheric sciences. We envision a data access system that will broaden the use of the ICESat dataset well beyond its core cryosphere community, will be ready to serve the upcoming ICESat-2 mission when its datasets come online in 2017~2018, and will provide pathways for the inclusion of similar datasets from other profiling altimetry missions.

Gao Chen/Langley Research Center
Subsetting Tools for Advanced Analysis of Airborne Chemistry Data (STA3CD)

15-ACCESS15-0025

Aircraft based measurements of atmospheric trace gas and aerosols represent a major component of NASA's Earth Observing System. Even in the era of satellites, airborne measurements remain a vital component due to the high spatial and temporal resolutions and the extensive variety of simultaneous measurements. Given these attributes, airborne observational data are well suited for characterization of atmospheric processes. In fact, the airborne data are often used as a benchmark for assessments and evaluation of both models and satellite observations, especially for the vertical distribution of atmospheric trace species. The Atmospheric Science Data Center (ASDC) developed and released the Toolsets for Airborne Data (TAD), a suite of web-based tools aimed to enhance airborne data discovery and usability, in September 2014. Thus far, TAD has a variety of users, including both scientists and researchers, spanning numerous agencies and universities, and has been used to create and download thousands of files. The current active users include PIs supported through the NASA ROSES 2013 solicitation entitled "Atmospheric Composition Campaign Data Analysis and Modeling (ACCDAM)" and researchers in the atmospheric modeling communities, e.g. AEROCOM and CCMI. While users have found the current TAD tools useful and effective, they have also requested certain feature enhancements. One of these requests is data subsetting for the in-situ atmospheric composition measurement data, which is typically the very first step in analyses involving airborne observations. To address the user community research needs, we propose to develop data subsetting tools as an addition to TAD. The proposed objective will allow users to extract data two ways:

1. Geographical: Users will be able to select data from a specific region across all relevant airborne sampling missions and flights.
2. Vertical Profile: Users will be able to identify individual aircraft vertical profiles of interest and obtain data associated with these vertical profiles.

The proposed activities will directly serve the Atmospheric Composition focus area by substantially improving the accessibility and usability of airborne data. Specifically, these tools will help a user more efficiently access the data relevant to their research and streamline the data analysis process. The geographical subsetter is expected to be frequently used in incorporating airborne data sets into model assessments and evaluations, while satellite validation would require coincident individual vertical profile data.

As a secondary objective we propose to develop the capability for TAD to output data in both ICARTT (a NASA ESDIS endorsed standard since 2010 for airborne measurement data) and netCDF file formats, which will make the data accessible to a wider user community, especially the modeling community.

All proposed objectives will be an extension of the existing TAD suite of tools, currently operating in a production environment at a TRL of 8. One major advantage of TAD is its extensibility. The system has been designed to allow new missions to be added, with only minor pre-processing, as they become available, and the proposed objectives will take

advantage of this. TAD includes extensive public documentation as well as a YouTube video tutorial that provides a complete start to finish user walkthrough. Both capabilities will be updated to include the proposed objectives, bringing TAD to an RRL of 6.

Members of this proposal team have been participants in multiple ESDSWG working groups, e.g. Airborne Data, ASCII, Open Source, and Geospatial Working Groups. In addition, the proposal team has been approved to chair an ICARTT Reformat Working group, which will work to update the ICARTT file format in order to enhance the usability, interoperability, and discovery of airborne data sets.

Thomas Clune/Goddard Space Flight Cente
NASA Open-Access Geo-Gridding Infrastructure (NOGGIn): An Integrated Service for Next-Generation Modeling, Analysis, and Retrieval Evaluation with Error Estimates
15-ACCESS15-0018

Perhaps the most common challenge to the systematic and routine comparisons of data from the model output/analysis with NASA remote sensing data is that of co-locating diverse datasets, often obtained from instruments with different observation characteristics, for the purposes of detailed comparison, e.g. Obs4MIPS (Teixeira et al 2014), and use, e.g. adaptation to climate change. Indeed, while a myriad of tools exist to locate, access, and visualize observational data, inter-comparison of disparate data sources requires tedious and often computationally intensive gridding or regridding, which is generally implemented in an ad hoc manner by individual users and becomes an obstacle to robust comparisons.

In addition to the duplication of effort stemming from the lack of standard gridding tools, ex-pediency on the part of individual researchers often results in sub-optimal gridded products that (1) limit accuracy due to failure to adequately account for spatial/temporal sampling bias, (2) lack robust gridding uncertainty estimates, and (3) omit the provenance, all of which limit the value of the products to other researchers. Simplicity and familiarity also drive researchers to apply traditional latitude-longitude (lat-lon) grids rather than better alternatives.

Moreover, the trend in the modeling community is to transition to next-generation grid sys-tems, such as geodesic and cubed-sphere, that possess superior quasi-equiareal, scalable charac-teristics. To maintain numerical stability, the CFL criterion for the time step used in the numeri-cal integration of the general circulation models must decrease with the smallest distance be-tween any two grid points. Since the meridians converge to a single point at each pole, as the model's spatial resolution increases, the smallest grid distance of a lat-lon grid system quickly approaches zero. Considerable computation will thus be wasted for high-resolution models based on lat-lon grids, for the short time step required near the poles becomes, progressively, a greater overkill for grid cells toward the equator. Approaches to filter the high-frequency signal to sidestep this concern introduce severe constraints on parallel performance.

Anticipating the need for converting and adapting NASA Earth science remote sensing data for compatibility with results from these next-generation models, we are proposing NOGGIn as an open-access service to enable routine and systematic gridding, co-location, and comparison of remote sensing data that not only makes adapting observations to these grids easy but also ad-dresses a number of gridding issues that currently plague researchers. In particular, this service will:

- Grid data onto next-generation icosahedral and cubed-sphere grids, as well as traditional lat-lon grids, with flexible temporal binning (e.g., by local time instead of UTC);
- Employ multiple estimation techniques for computing grid-cell values, including kriging for reducing sampling bias;
- Provide robust uncertainty estimates of gridded values;
- Embed provenance meta-data, describing the operations applied, as well as the original data sources, to improve traceability; and
- Allow for introduction of additional estimation methods with a modular design.

We will partner with MODIS Adaptive Processing System (MODAPS) to leverage and aug-ment its existing map projection web services and to eventually extend NOGGIn as a service in the cloud. We will develop a web client that provides a user interface to these services, making them accessible through a web browser.

Our Earth Science Focus Area is “Water and Energy Cycle”. We will design and implement NOGGin to “improve data access, management and interoperability” with the intent to “increase the efficiency for the user and enable new users to benefit Earth” and be compatible with NASA’s “distributed heterogeneous data and information system architecture” and NASA’s Earth Science “system of systems” infrastructure.

Larry Di Girolamo/University of Illinois at Urbana-Champaign
ACCESS to Terra Data Fusion Products
15-ACCESS15-0043

Terra is the flagship of NASA’s Earth Observing System. Launched in 1999, Terra’s five instruments continue to gather data that enable scientists to address fundamental questions that are central to the six NASA Earth Science Research Focus Areas. It is amongst the most popular NASA datasets, serving not only the scientific community, but also governmental, commercial, and educational communities.
The strength of the Terra mission has always been rooted in its five instruments and the ability to fuse the instrument data together for obtaining greater quality of information for Earth Science compared to individual instruments alone. As the data volume grows and the central Earth Science questions shift from process-oriented to climate-oriented questions, the need for data fusion and the ability for scientists to perform large-scale analytics with long records have never been greater. The challenge is particularly acute for Terra, given its growing volume of data (> 1 petabyte), the storage of different instrument data at different archive centers, the different file formats and projection systems employed for different instrument data, and the inadequate cyberinfrastructure

for scientists to access and process whole-mission fusion data (including Level 1 data). Sharing newly derived Terra products with the rest of the world also poses challenges. This proposal, ACCESS to Terra Data Fusion Products, aims to resolve two long-standing problems:

1) How do we efficiently generate and deliver Terra data fusion products?
2) How do we facilitate the use of Terra data fusion products by the community in generating new products and knowledge through national computing facilities, and disseminate these new products and knowledge through national data sharing services?

Our approach leverages national facilities and services that are managed by the National Center for Supercomputing Applications (NCSA), specifically the National Petascale Computing Facility, which houses the Blue Waters supercomputer, and the National Data Service (NDS). Key advantages of leveraging Blue Waters and the NDS for access, usage, and distribution of Terra data fusion products and science results are that the Terra data and processing are local, with access and sharing that are global. This represents a significant community-element addition to NASA's system of systems infrastructure. ACCESS to Terra Data Fusion Products will initiate the development, access and delivery of Level 1B radiance Terra Fusion files for the broader community. Level 1B fusion provides the necessary stepping-stone for developing higher-level products, and provides the framework for other flavors of fusion. Enhancements to our existing open source codes in the CyberGIS Toolkit for scalable map projections on any grid for the new Terra Fusion files will also be delivered.

The potential impacts here are broad. In particular, this project will allow scientists and other Terra users to (1) facilitate greater ease in creating new geophysical retrieval algorithms that provide greater accuracy relative to the current single instrument algorithms, (2) provide an easy mechanism for users to access and process the entire Terra record, (3) greatly reduce error and redundancy in the science community who are using multiple Terra instrument datasets, (4) provide greater insight into geophysical processes through synergistic use of different instruments and their products, and (5) provide a framework for fusion that could extend to other NASA missions and constellations (e.g., AM-constellation, DSCOVR, GOES-R, A-Train). The end result will be to facilitate discovery and to accelerate progress in Earth Science research.
To reach our goals, we bring together expertise from NCSA, HDF Group, and the University of Illinois' Department of Atmospheric Sciences, with collaborators from NASA centers and other institutions, to form a lasting working relationship with NASA ESDIS.

Andrea Donnellan/Jet Propulsion Laboratory
Enhanced Access to NASA Heterogeneous Geodetic Imaging Products through GeoGateway Analysis, Modeling, and Response Tools
15-ACCESS15-0009

GeoGateway's enhanced access to NASA geodetic imaging data will steer users to relevant products using automated feature extraction, manual search tools, and data

quality metrics. The objective of GeoGateway is to allow users to efficiently find and use NASA geodetic imaging data products. We will integrate new functionality to access, analyze, and model GPS and UAVSAR data, with incorporation of spaceborne InSAR, Landsat, and optical imagery products as appropriate. GeoGatway goals are to to a) simplify the discovery of geodetic imaging products; b) enable researchers to explore and integrate data products through online or offline analysis; c) allow researchers to easily share, publish, and collaborate on results of these online experiments. We will increase the value of existing geodetic imaging products to researchers through automated machine learning and computer vision techniques to identify features, anomalies, and artifacts. The objective of this project is to bridge the gap between production and end-use of data products allowing users to efficiently find and use NASA geodetic imaging data products and integrate them with other data products. The data products are far more powerful when they are used in conjunction with other products to test a hypothesis. Automated and manual approaches will be available for the user through the interface. Users will be able to perform time and geographically limited search for products. Time windowing is particularly important because crustal deformation velocities vary over time as earthquakes and postseismic deformation occur, causing offsets and rate changes. Users will be able to plot results by different reference stations highlighting different tectonic characteristics. We will enhance our rating and annotation system so that users can tag products by quality. User ratings will be aggregated and available for search and sorting. Automated feature detection will identify edges, discontinuities, anomalies, and data product quality. Statistics on the GPS and UAVSAR data will enable the user to separate data artifacts and errors from geophysical signals. This project is directly relevant to the ACCESS solicitation because it meets the primary objective, which is to enhance, extend, and improve existing components of NASA's distributed and heterogeneous data and information systems infrastructure. NASA heterogeneous crustal deformation data come in the form of GPS position time series and radar or optical maps of components of surface deformation. Earthquakes and postseismic response complicate the already heterogeneous data by adding episodic jumps and changing rates of surface deformation to the products. Data volumes will become even larger and the surface deformation data more complicated when NASA launches NISAR. This forward-looking project will use existing data as a testbed for the NISAR mission. GeoGateway components are open source and available for reuse by other projects and researchers. The emphasis on facilitating routine workflow promotes the reuse of techniques necessary to advance earthquake research. GeoGateway contains the infrastructure to access distributed data and tools to improve data usability.

Riley Duren/Jet Propulsion Laboratory
Methane Source Finder
15-ACCESS15-0034

The Methane Source Finder (MSF) project will implement a science data system (SDS) capable of streamlined processing and visualization of airborne methane remote sensing data products. NASA imaging spectrometers AVIRIS-NG and HyTES are effective at surveying large areas and generating high-resolution maps (1-3 meters per pixel) of atmospheric methane enhancements in the aircraft subcolumns. Surveys using both

sensors in 2013-2015 have successfully detected hundreds of methane point sources during surveys in California, New Mexico, and Colorado. There is strong interest in methane point source detection and attribution by stakeholders in government agencies and private sector. These data offer great value to these stakeholders if these methane data products are generated with low-latency and accessible by non-expert users. The current manual analysis workflow is very time-consuming, requires specialized expertise and could greatly benefit from increased automation and integration. The MSF SDS will be capable of rapid generation of L3 products via the MSF data pipeline subsystem. These L3 products will then be ingested by a analysis and visualization subsystem. The MSF system will leverage and integrate several existing NASA data system tools. The system will create a L4 data output file that helps users characterize and confirm methane point sources.

We are collaborating on the MSF project with two key State of California stakeholders: (a) the California Air Resources Board (CARB), and (b) the California State Energy Commission (CEC). Our second project objective is to work with these collaborators to demonstrate the use of the MSF Science Data System under operational conditions during a planned airborne methane survey over the state of California in 2016.

John Readey/The HDF Group
Object Store-Based Data Service for Earth System Science
15-ACCESS15-0031

Cloud computing offers key benefits of scalability, cost, and redundancy compared with traditional data centers. However most of the tools and software systems developed for NASA data repositories were not developed with the cloud in mind and do not fully take advantage of cloud-based technologies. To address this, we propose the development of a web service (Highly Scalable Data Service or HSDS) that will use object storage rather than file based storage to vend earth science data. Object storage services are provided through all the leading public (Amazon AWS, Microsoft Azure, Google Cloud, etc.) and private clouds (Open Stack), and provide the most cost-effective means of storing large data collections online. In addition to being cost-effective, the service will be highly scalable in that it will be able to support an effectively unlimited number of clients at high throughput. To enable compatibility with existing tools and applications, we will create client libraries that are API compatible with existing libraries for HDF5 and NetCDF4. We will demonstrate the capabilities of this service by setting up a platform on Amazon AWS that integrates the service with existing NEX/OpenNEX datasets and tools.

Vardis Tsontos/NASA Jet Propulsion Laboratory
Leveraging Available Technologies for Improved Interoperability and Visualization of Remote Sensing and In-Situ Oceanographic Data at the PO.DAAC
15-ACCESS15-0017

The earth science enterprise relies on the integration of multivariate data from diverse observational platforms. Whether for NASA mission cal/val, science or decision support, the coupling of remote sensing and in-situ field data is integral to oceanographic workflows. This need is reflected by the PO.DAAC, historically having a remote sensing focus, now adapting to additionally support NASA field campaign datasets. However, the inherent heterogeneity of in-situ datasets and their general lack of adherence to meta/data standards is a significant impediment to interoperability. The problem originates early in the data lifecycle with instrument manufacturer software typically not supporting output to self-describing data files with standards compliant metadata and propagates, impacting long-term stewardship and usability of data. Post-hoc conversions are complicated, thus substantially increasing the cost of archival at DAACs and utilization in end-user applications.
This proposal aims to leverage existing technologies to address the aforementioned technical challenges relating to in-situ/satellite data interoperability, with a focus on a representative class of oceanographic field data: data from electronic tags deployed on a variety of marine species. Fueled by the affordability of the hardware and symptomatic of the importance of the data, tagging applications are a rapidly proliferating area for applied fisheries management and oceanographic studies. There are several large scale tagging projects and operational tagging programs run by fisheries management agencies

(eg. NOAA/NMFS) providing extensive georeferenced, depth resolved time series of environmental data. These datasets are plagued, however, by the data standardization issues mentioned. Consequently, few of these many, invaluable datasets are accessible via national data archives despite requirements, with significant implications for their long-term stewardship.
The proposed work involves the enhancement and integration of available high TRL components within PO.DAAC for improved interoperability and support of in-situ data: Unidata's Rosetta tool, NODC's netCDF in-situ templates, PO.DAAC's DMAS system, the Tagbase relational database system, a comprehensive data visualization framework developed at JPL (CMC), and THREDDS, extended here for improved support of point/trajectory spatial data types. The resulting system will be tested in the context of a range of in-situ (tagging) and PO.DAAC satellite data holdings. We intend to demonstrate, deliver and ultimately sustain operationally a reusable and accessible set of tools to: 1) mediate reconciliation of heterogeneous source data into a tractable number of standardized formats; 2) harmonize existing metadata models for satellite and field datasets within DMAS; 3) demonstrate the value added of integrated data access via a range of available tools and services hosted at the PO.DAAC, including a web-based visualization tool for integrated mapping of satellite and in-situ data.
We also recognize the importance of tackling data standardization issues at source. An innovative part of our project plan involves partnering with the leading electronic tag manufacturer (Wildlife Computers) to promote the adoption of appropriate data standards in their processing software. The proposed project thus adopts a model lifecycle approach complimented by broadly applicable technologies to address key data management and interoperability issues for in-situ data. Outcomes will: 1) enhance NASA capacity in the stewardship of a wide range of earth science data consistent with user needs and NASA earth/applied sciences and national IOOS (Integrated Ocean Observing System) priority areas. 2) Address critical constraints to in-situ data interoperability, archival and access per federal agency mandates and most recently under the Big Earth Data Initiative (BEDI).

Petr Votava/Ames Research Center
(Open)NEX: Enabling Code-to-Data Migration between High-Performance Computing, Cloud and Beyond
15-ACCESS15-0035

It is becoming increasingly apparent that moving code to data is a more efficient way to conduct research in the Earth sciences. However, in many instances scientists still prefer to move data because science codes tend to be hard to port to different systems. In this proposal we address the code-to-data migration by integrating the NASA Earth Exchange with Docker, a popular open-source Linux tool. This will enhance deployment and discovery of NASA tools and data services and greatly increase software reuse.