A STANDARDS-BASED DYNAMIC DISTRIBUTED DATA INFRASTRUCTURE FOR EARTH SCIENCES

STEFAN ACHLEITNER[5], WERNER BENGER[1], GERD HEBER[4], HARTMUT KAISER[1,2], NAYONG KIM[1], MARCEL RITTER[5], JIAN TAO[1], AND SHANTENU JHA[3,1]

1. Challenges for Dynamic Distributed Data in Earth Sciences

Scientific progress in the Earth Sciences is predicated on the ability to share, analyze and archive extreme amounts of data, which are often collected dynamically from a widely distributed sources. Sharing, analyzing and archiving these data is a very important but also challenging task facing the Earth Science community. Providing an integrated compute and data cyberinfrastructure that can support application requirements in a scalable, extensible and interoperable fashion is the primary challenge.

In this white paper, we outline several use cases and discuss some characteristics and approaches towards a solution. An important requirement in the process is to create a unified data model and an open management system for dynamic distributed data. Specifically, we propose developing and deploying a standards-based multi-layered infrastructure for dynamic distributed data to (1) compute, manage and integrate dynamic distributed data, and (2) unify data model. While we promote standards, we emphasize the adoption of existing technology and software in the public domain.

The multi-layered structure of the proposed integrated compute-data cyberinfrastructure provides the right abstractions that hides implementation details that prevent scalability, extensibility and interoperability. We introduce an Earth Science Toolkit layer, as a placeholder for a community need-driven software package providing required tools to leverage the increasingly sophisticated cyberinfrastructure without "vendor lock-in". As we will discuss, our proposed integrated compute-data infrastructure development addresses the requirements of multiple application scenarios and communities.

2. A Standards-Based Cyberinfrastructure

2.1. **Standards-based Cyberinfrastructure: A Moral and Technical Imperative.** There is pressure on modern cyberinfrastructure to support more science with *less* – hardware, software and personnel, at every level. To address the later two categories, there is critical need to share best-practices, reduce software development redundancy, and take a systems-approach to designing cyberinfrastructure. Standards promote all three. In addition to the above "soft issues", there are other specific tangible advantages to using standards in the design of cyberinfrastructure: we focus on interoperability and extensibility (modular)

We focus on open standards so as to ensure extensibility across a broad spectrum of the applications requirements and scenarios. The use of open standards also reduces risk, maximizes return on investment, and future-proofs critical applications. It increases the ease and likelihood of interoperability between all involved parties.

(1) Center for Computation & Technology, Louisiana State University, Baton Rouge, LA, USA.
(2) Department of Computer Science, Louisiana State University, Baton Rouge, LA, USA.
(3) Rutgers University, USA.
(4) The HDF Group, IL, USA.
(5) Unit of Hydraulic Engineering, Department of Infrastructure, Innsbruck, Austria.

Given the complexity of developing and provisioning production-grade and scalable cyberinfrastructure [1] for modern scientific applications, the ability to take a systems-level approach to CI development is mandatory. A systems-level approach in turn demands modular and componentized development, and in order to minimize incompatible components/modules, a standards-based development regimen is required.

Standards are a promising and secure way to achieve different levels of interoperability – both system-level and application-level is required in order to facilitate easy-access to diverse infrastructure, as well as as a basis for scalability and extensibility. In particular, standardized interfaces to functionalities, data, or services ensure that different parts of a system can be seamlessly integrated. In the past it was a "desirable feature" while in today's science it is a "necessary feature".

In other words, the use of effective standards – where they exist, is not just a technical nicety but a moral imperative. The imperative is made more important when trying to build a community infrastructure and effort.

In this white paper, we propose a technology solution that uses some of the open standards or "de facto" standards for the EarthCube data infrastructure. We briefly outline the existing standards – *de jure* and *de facto*, that we believe can be used to build cyberinfrastructure for the Earth Science community.

The Open Grid Forum (OGF [9]**)** is an open community committed to driving the rapid evolution and adoption of applied distributed computing. OGF accomplishes its work through open forums that build the community, explore trends, share best practices and consolidate these best practices into standards. Many of the standards developed by OGF have been widely adopted in the community and have proven to foster interoperability in large scale projects.

The HDF Group [11] provides a unique suite of technologies and supporting services that make possible the management of large and complex data collections. The Hierarchical Data Format V.5 (HDF5), developed by the HDF Group is a "de facto" standard format for storing scientific data. The NetCDF-4 [14] format that is widely used in the Earth Science community is also built upon HDF5. Being a member of **the Open Geospatial Consortium (OGC** [15]**)**, the HDF group will be collaborating on OGC initiatives, and develop relationships with other consortium members.

When implementing Earth Science applications upon the proposed data infrastructure, developers may create toolkits based on some community standards, such as **OGC** that provides publicly available interface standards which support interoperable solutions that "geo-enable" the Web, wireless and location-based services and mainstream IT.

We believe a successful data infrastructure based on these open standards will not only directly benefit the Earth Science community with an open and sustainable data infrastructure, but also help to propagate and improve those standards. The incorporation of existing frameworks, such as GDAL [10], OGR [16] or similar, which show promising success in terms of providing uniform data abstraction and are widely used, will ensure community wide acceptance and usability.

3. Design, Development and Integration

In order to capture user requirements, our design process analyzed several use-cases. **The Coastal Emergency Risks Assessment (CERA) group** [6] is a coastal modeling research & development effort providing operational advisory services related to impending hurricane events and other coastal hazards. CERA has started dealing with different data formats for years. The success of CERA lies in the fact that CERA employed a standards-based approach, namely, CERA used the OGC standards to discover, encode, and represent data using OGC interfaces and OGC conformant tools. In the white paper by Chen et al. [12] on tsunami simulations, the authors present **an earthquake-generated tsunamis scenario**, a workflow management system leveraging a distributed computing environment becomes a very important component in the project. However, to ensure seamless data exchange and automatic data driven computing, standards come in naturally to enable such an integration.

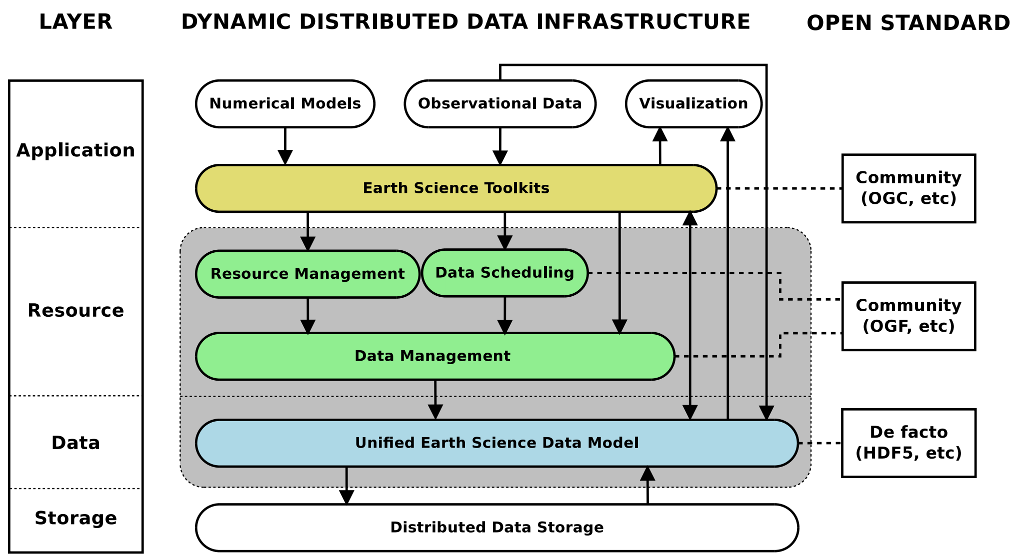

FIGURE 1. A conceptual overview of a standards-based dynamic distributed data infrastructure for community-based collaborative research in Earth Sciences.

The Ensemble Kalman Filter (EnKF) is the basis for many Earth Science applications. We have implemented [8] an standards-based EnKF workflow, that manages the distribution of computational tasks; this can be used by multiple applications ranging from reservoir simulations to climate prediction. We present EnKF as a more detailed use-case later in this paper.

Based on our experiences in various large scale collaborative projects, we envision a 4-layer structure (see figure 1) of a standards-based dynamic distributed data infrastructure for community-based collaborative research in Earth Sciences. The 4 layers are *application*, *resource*, *data*, and *storage*. Although the Earth Science observations and models together with the data storage systems are crucial components of the overall EarthCube data infrastructure, they are not the primary concerns of this white paper, in which we focus on the data and resource layers that serve as a bridge between Earth Sciences and hardware facilities.

3.1. **Resource Layer.** From an analysis [2, 3] of many applications – some of which are in the Earth Sciences Domain, whilst others have the general application characteristics akin to the the dynamic data problems, we find that there is a need for middleware and cyberinfrastructure solutions that: (R1) provide first-class support for the coupling of distributed data – both static and dynamic, to computation using abstractions and middleware constructs that go beyond the current state-of-the-art, (R2) support performance trade-offs across distinct levels, and (R3) enable integrated end-to-end capabilities.

SAGA [17] is a high-level programming interface that provides the basic capabilities for developing loosely-coupled, multi-component distributed — physically and logically — applications and services. Applications use the SAGA API to access platform specific implementations of a desired functionality. SAGA is also serves as middleware, as it is engineered to provide the integrative tissue between the application and the system. SAGA has been successfully demonstrated to support compute-driven dynamic workloads, similar to those that arise in uncertainty quantification, on high-end machines (e.g., on 30,000 cores on Kraken at NICS/ORNL).

3.2. **Data Layer.** Interoperability across Earth Science applications is a recognized need, as comparing data from different sources is a core part of Earth Science research. The use of distinct file formats for the same type of data due to optimizations and specializations for each specific case, or even incompatibilities of the used data models itself [13], is a hurdle for availing and sharing data beyond their original creation domain, both with respect to applicability in another context, for instance visualization [5], as well as for archival since knowledge about a specific data layout may be lost over time. The same problem occurs in other sciences, such as biomedical imaging [7]. It is thus an general problem across disciplines, and a discipline-independent solution will be most future-safe. To provide the widest applicability and interoperability a solution should tackle the problem on different specialization levels: (1) aspects of the dataset(s) which are purely mathematical properties (for instance, the dimensions and numerical precision of a dataset) and can thus be understood most widely, (2) aspects of the dataset(s) which are specific to Earth Sciences (for instance the coordinate system used), but are generally understood and agreed on within this domain, (3) application-specific aspects which are important within one very narrow usage scenario, but are not relevant or transportable to other applications.

A solution, such as a file format or database access method, allowing to clearly identify and store these different kinds of data properties within the same data entity will ensure the widest interoperability and "sharability" of datasets. Thereby, a dataset can be interpreted even beyond Earth Sciences just based on its mathematical properties, allowing to apply (for instance) biomedical image processing methods on Earth Science images without any modification of applications. Datasets exhibiting scientific metadata can be understood within the Earth Science community while application-specific extensions will still be sustained without loss of information in such a data sharing scenario.

A very promising candidate offering all the features to implement this ideal scenario is HDF5, an domain-agnostic high performance I/O library and file format. It is well tested in the domain of high performance computing since more than a decade. HDF5 is actively managed by the HDF group [11], and in many application domains, HDF5 as well as NetCDF-4, which has HDF5 under the hood, have become de-facto standards in Earth Sciences. Due to its flexible, hierarchical structure it allows to specify different kinds of dataset properties.

3.3. **Ensemble Kalman Filter - A Case Study.** A typical Ensemble Kalman Filter (EnKF) workflow is shown in figure (2a) [4]. The EnKF uses an ensemble of models to estimate the error statistics of the model with observational data. The EnKF has its applications in various fields in Earth Sciences, such as ocean and atmospheric modeling. As shown in figure (2a), an EnKF workflow iterates through 5 steps: (1) Sensor data is steamed to the data spool. (2) Available Sensor data is queried by the workflow manager. (3) Another ensemble of simulation is launched. (3-1) Alternately "sleeps" and check the spool again by workflow management system. (4) Simulation data is collected EnKF is launched. (5) Control returned to Standards-based workflow management system, workflow restart at stage(1).

A conceptual diagram for a typical EnKF workflow in the multi-layered dynamic distributed data infrastructure can be found in figure (2b). The sensor data is now saved on the distributed storage in a unified data model. The workflow management system will take the data driven workflow provided by the user and co-schedule both workflow and data flow. Besides easy data sharing and utilization due to the unified data model, the major advantages of the scenario shown in figure (2b) are: data are stored distributed to incur less load on a central server; workflow is separated from the application to enable quick switch to other applications.

4. Relation to Community Efforts, Collaboration and Broader Impact

The proposed standard-based multi-layered dynamic data infrastructure serves as a common bridge to exchange data across the Earth Science community. It will lead to an enhanced communication potential of academic partnerships on a national as well as international level. A common agreement on a unified data platform will be essential for easing scientific discovery in Earth Science and beyond. A large scale community level collaboration is necessary to make it happen. The success of the data infrastructure in EarthCube could potentially be copied to other communities by switching the application layer. In multidisciplinary researches involving distantly related fields, extra APIs could be defined. The sharing, analyzing, and utilizing data of distantly related fields can be achieved with minimal changes in the application codes.

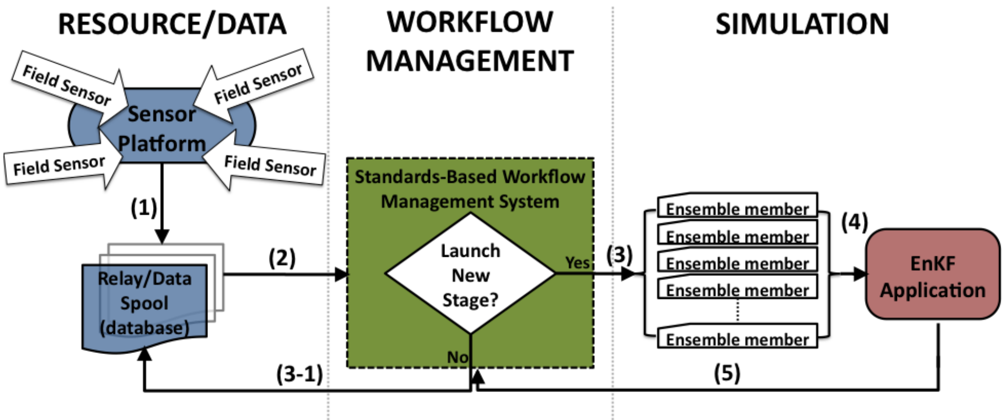

(a) A typical EnKF workflow

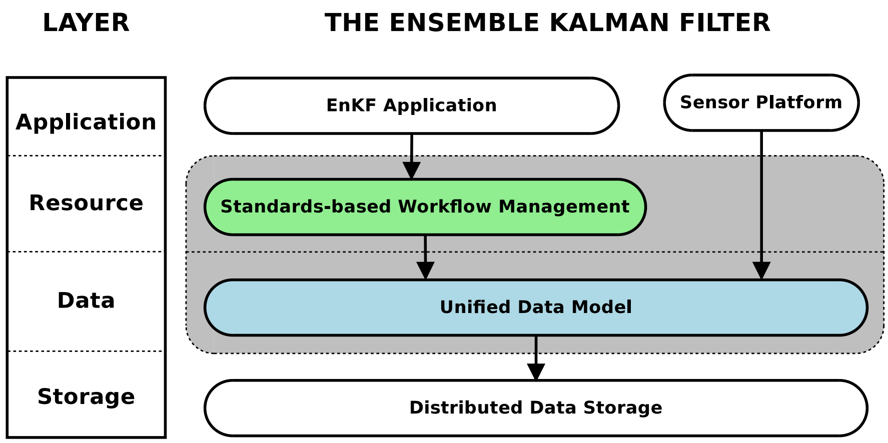

(b) An EnKF workflow in the multi-layered data infrastucture

FIGURE 2. Figure (a) shows the control flow of typical end-to-end compute-data integration scenario for EnKF (adapted from Ref. [8]. Figure (b) shows a conceptual diagram for a typical EnKF workflow in the multi-layered dynamic distributed data infrastructure.

5. Future Plans

We will continue our efforts to take the standards-based approach to address the challenges in dealing dynamic distributed data in Earth Sciences and beyond. We will adopt and promote widely used open standards – *de jure* and *de facto*, such as OGF, OGC and HDF5, in our development of the data infrastructure. There will be 5 milestones that we want to achieve: (M1) first implementation of a dynamic distributed data management system; (M2) working with Earth scientists and engineers to release an Earth Science Toolkits to provide common APIs to link existing Earth Science applications to the data infrastructure; (M3) release a stable version of the dynamic distributed data management system and the Earth Science Toolkits; (M4) refine the data infrastructure and initiate some community wide standardization process for a unified data model and APIs for the data infrastructure; (M5) first implementation of EarthCube.

References

[1] D. S. Katz, S. Jha, M. Parashar, O. Rana, and J. Weissman. Survey and Analysis of Production Distributed Computing Infrastructures. Technical Report CI-TR-7-0811. Computation Institute, University of Chicago & Argonne National Laboratory. http://www.ci.uchicago.edu/research/papers/ CI-TR-7-0811.

[2] S Jha *et al*, White Paper Submitted to DOE ASCR (Rich Carlson), August 2011 http://www.cct.lsu.edu/~sjha/select_publications/D3-exa_jha_katz_weissman.pdf.

[3] S Jha *et al*, Distributed Dynamic Data-intensive Programming Abstractions and Systems http://www.cct.lsu.edu/~sjha/select_publications/3dpas.pdf.

[4] Developing Autonomic Distributed Scientific Applications: A Case Study From History Matching Using Ensemble Kalman-Filters, In Proceedings of the IEEE 6th ICAC, Workshop on Grids Meets Autonomic Computing (Barcelona, Spain, June 15 - 15, 2009) http://www.cct.lsu.edu/~sjha/dpa_publications/lazarus_gmac09.pdf.

[5] Werner Benger. On safari in the file format djungle - why can't you visualize my data? *Computing in Science & Engineering*, 11(6):98–102, 2009. **Feature Article in "Computing Now"** http://www.computer.org/portal/web/computingnow/1109/whatsnew/cise.

[6] CERA - Coastal Emergency Risk Assessment. http://nc-cera.renci.org/.

[7] Matthew T. Dougherty, Michael J. Folk, Herbert J. Bernstein, Frances C. Bernstein, Kevin W. Eliceiri, Werner Benger, Erez Zadok, and Christoph Best. Unifying biological image formats with hdf5. *Communications of the ACM*, 52(10):42–47, 2009.

[8] Yaakoub el Khamra, Shantenu Jha, and Christopher D. White. Modelling data-driven co2 sequestration using distributed hpc cyberinfrastructure. In *Proceedings of the 2010 TeraGrid Conference*, TG '10, pages 6:1–6:8, New York, NY, USA, 2010. ACM.

[9] Open Grid Forum.

[10] GDAL - Geospatial Data Abstraction Library (GDAL). http://gdal.org/.

[11] The Hierarchical Data Format (HDF) Group. http://www.hdfgroup.org.

[12] Integration of Tsunami Source Database and Tsunami Generation, Propagation and Coastal Impact Model for GeoHazard Prediction and Mitigation. EarthCube White Paper.

[13] S. Nativi, B. Blumenthal, T. Habermann, D. Hertzmann, R. Raskin, J. Caron, B. Domenico, Y. Ho, and J. Weber. Differences among the data models used by the geographic information systems and atmospheric science communities. In *Proceedings American Meteorological Society - 20th Interactive Image Processing Systems Conference*, 2004.

[14] NetCDF (network Common Data Form). http://www.unidata.ucar.edu/software/netcdf.

[15] Open GeoSpatial Forum (OGC). http://www.opengeospatial.org/.

[16] OGR Simple Feature Library. http://gdal.org/ogr/index.html.

[17] The SAGA Project. http://www.saga-project.org.