Extreme Scale Data Management, Analysis, Visualization, and Productivity in Climate Change Science

Panel Report

Dean Williams, Lawrence Livermore National Laboratory
Don Middleton, National Center for Atmospheric Research
Mihai Anitescu, Argonne National Laboratory
Venkatramani Balaji, Geophysical Fluid Dynamics Laboratory
Wes Bethel, Lawrence Berkeley National Laboratory
Steve Cotter, Lawrence Berkeley National Laboratory
Gary Strand, National Center for Atmospheric Research
Karen Schuchardt, Pacific Northwest National Laboratory
Arie Shoshani, Lawrence Berkeley National Laboratory

Envisioning Today the Extremes of Tomorrow: Toward Strategic Action on the Climate Change Imperative

Current Status

Introduction

Projections of potentially catastrophic changes to human habitat associated with anticipated climate change make it imperative to develop a "built-to-share" scientific discovery infrastructure that is funded nationally and internally for the benefit of climate researchers, policymakers, and society. Currently, climate researchers have a difficult time locating and analyzing data for scientific studies, which leads to untenable delays and gaps in the evolving scientific understanding of our changing world. Existing data sharing processes associated with data discovery, access, analysis, and visualization process are labor intensive and collaborative communication and governance mechanisms are inadequate. Today, human intervention is needed at every phase of data management, access, analysis and visualization.

Building on the Earth System Grid

Making available extreme-scale data, distance computing, and distributed applications from the network for participating national and international climate research institutions requires an infrastructure that is realized in embryonic form as the Earth System Grid (ESG). Moreover, to improve research ability and productivity and to solve complex scientific climate problems, this next-generation infrastructure requires advances in all sub-components of the ESG science environment needed to support true distance computing, databases, visualization, and other distributed applications. Projecting a glimpse of the future, the ESG Center for Enabling Technologies is striving to (1) make data more accessible to climate researchers by making access location transparent to heterogeneous datasets; (2) meet specific needs of national and international climate projects for distributed database, data access, and data movement; (3) provide secure web-based access portals to multi-model data collections; and (4) provide a widerange of Grid-enabled climate data analysis tools and diagnostic methods to international climate centers and U.S. government agencies.

Vision

Future technological advances envisioned over the next 5-20 years must address extreme-scale data warehousing, scalability, and service-level requirements that far exceed what exists today. During this period of rapid expansion, users worldwide will have access to hundreds of exabytes of data, which will need to be stored at multiple disparate sites and connected to exaflop servers operating with high reliability at unprecedented Internet speeds. Scientists and policymakers will be able to retrieve information and derive sophisticated data products within milliseconds. Transcending geographical and organizational boundaries, scientists will employ virtual collaboration analysis tools to instantaneously manipulate data that are viewable by all users. These tools, unlike the current static analysis tools, will support the co-existence of many users in a productive shared virtual environment. This advanced technological world, driven by extreme-scale computing and the data it generates, will increase scientists' productivity, exploit their national and international relationships, and push their research to new levels of understanding.

Use Case (sidebar)

The following use case illustrates the interface of the researcher/user with the data in order either to perform scientific research or to understand environmental concerns relevant for setting policy. The processing effort involves moving vast amounts of data (spanning several government agency analysis centers) to and from various sites around the world.

In this example, domain experts (i.e. climate scientists) in different international locations are called upon to provide crucial inputs for policy decisions on outbreaks of African malaria that depend on local climate conditions. From their remote sites, the scientists search a climate portal containing exabytes of high-resolution regional central Africa data. When their data search proves unproductive, the scientists run several models in real-time to generate ensemble simulations of African climate. Using server-side visualization tools, they are able to simultaneously view and annotate plots of ensemble climate statistics on their respective platforms. The climate scientists then save this session, and later a malaria policymaker discovers (using a "new search capability") the provenance of this saved session. Working with the scientists, the policymaker conducts further assessment and re-analysis of the derived datasets before reducing to 20 terabytes (from their original size of 20 petabytes) and moving them to a local workstation for further study. The scientists then integrate the climate model ensemble data with the African malaria data for potential future use.

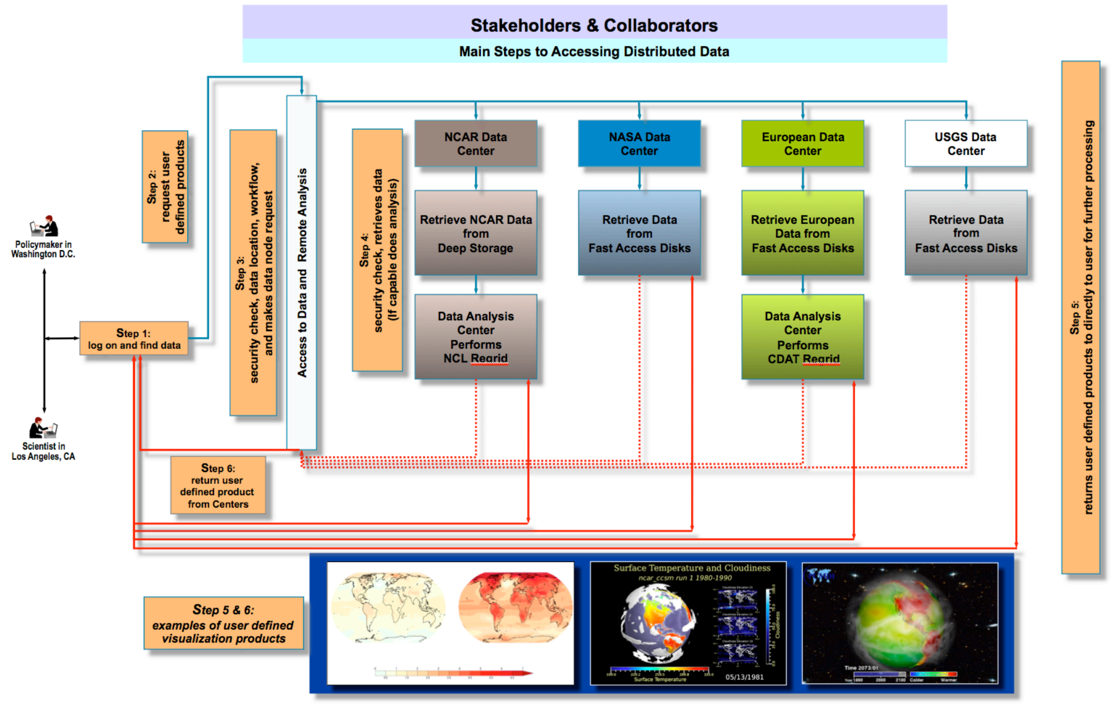

Figure 1. Schematic depiction of a use-case scenario supporting remote login for experts (e.g., model developers, climate researchers) and non-experts needing fault-tolerant end-to-end system integration and large data movement. For the expert user, the system expresses all the capabilities and complexities needed for rich data exploration and manipulation. For the non-expert, however, a simple abstraction layer includes easy-to-use controls for maneuvering within the system.

Thus, although the targeted primary users are domain experts, it is essential that non-experts (e.g., politicians, decision makers, health officials, etc.) also are able to access much of the data. Rich but simple interfaces will allow the non-experts to accomplish difficult data manipulation and visualization tasks without having to understand the complexities of application programs or the computing environments on which they depend (Figure 1).

Wide Area Network Issues and Challenges

Climate model datasets are growing at a faster rate than the dataset size for any other field of science (Figure 2). Based on current growth rates, these datasets will be hundreds of exabytes by 2020. To provide the international climate community with convenient access to these data and to maximize scientific productivity, these data will need to be replicated and cached at multiple locations around the globe. Unfortunately, establishing and managing a distributed data system presents several significant challenges not only to system architectures and application development, but also to the existing wide area and campus networking infrastructures. For example, transport technologies currently deployed in wide area networks do not cost-effectively scale to meet the scientific community's projected aggregate capacity requirements based on the growth rates for dataset size. Even if backbone network technology improvements increase link speeds from the current 10 Gigabits per second to 100 Gigabits per second and are in production service by 2012, as anticipated, more efficient use of networking resources will be essential. Efforts are underway to develop hybrid networks with dynamic virtual circuit capabilities, such as those networks currently being tested and deployed by research and development networks like ESnet. Although dynamic virtual circuits allow high-capacity links between storage and computer facilities to be created as needed and then deactivated quickly to free up network capacity for other users, much work is still required to optimize and harden the software.

Easy and efficient data transport across wide area networks is not the only networking challenge to be faced over the next 5-10 years. In the use case mentioned above, the policymaker easily downloaded 20 TB of data to their site. To achieve easy downloads and better performance for the ordinary user, training is key. Often there is too little staff or outreach for this effort—leaving the user of the system frustrated.

To maximize throughput between distributed systems, sophisticated network monitoring tools are required to enable real-time monitoring of the entire end-to-end network link—including campus networks—so as to avoid congestion and provide the necessary performance data to enable fine tuning of the protocol stack and assist with troubleshooting. This network monitoring software will need to be tightly integrated, preferably facilitated via the development of applicable standards, with the software systems that control data movement and dynamic circuit control to achieve our networking goals. Data movement infrastructure—both hardware and software—will have to be maintained and upgraded to enable the continued efficient use of the network. Long-term investments in middleware and test/measurement tools are critical. The implication is that tools will become more essential over time, and thus will need to be officially supported with long-term funding.

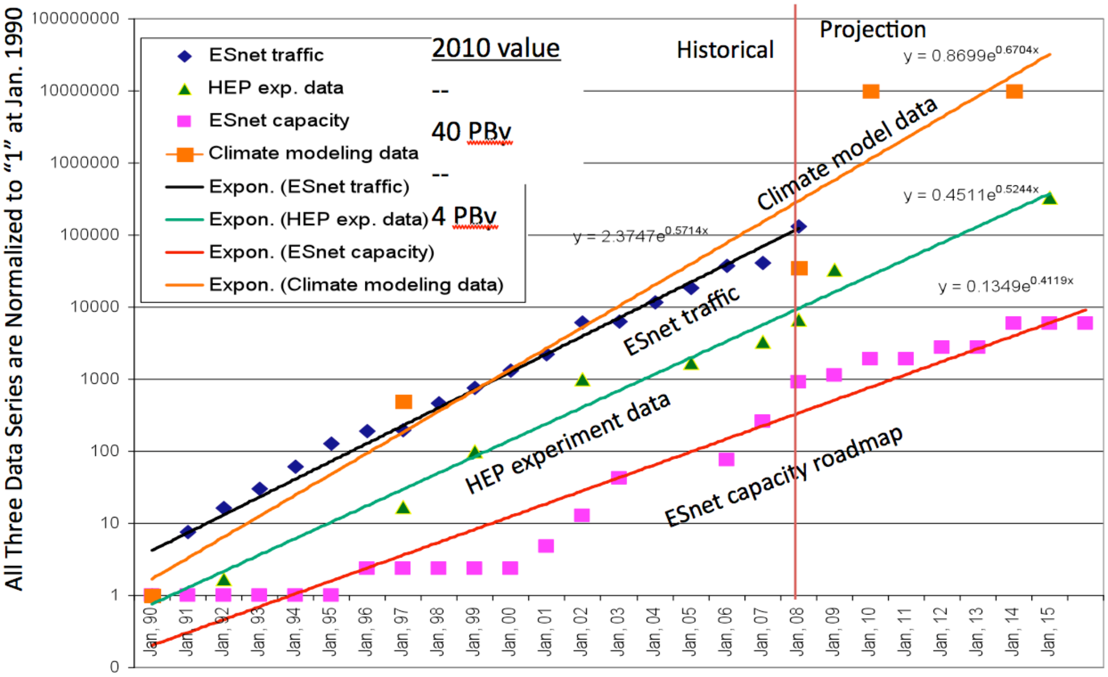

Figure 2. Long-term trends in future network traffic and capacity.

Data Management Issues and Challenges

Managing collections of extreme-scale scientific databases, consisting of diverse Earth system data sources, present several significant technological challenges. As mentioned previously, data archives will be distributed over a wide geographic area in federated databases. While such data archives will be diverse depending on the country or agencies involved, they should have uniform data access layers that permit them to operate transparently. Furthermore, adding new archives and data to the infrastructure should be a simple dynamic process that does not disrupt ongoing operations. Automating monitoring of storage usage and the effectiveness of data access will be essential to identify and correct bottlenecks. In moving toward increased collaboration among groups of researchers, the technical software infrastructure must support shared access to distributed data archives and applications, and must provide flexible software systems designed for individual or group needs. Fundamental to this infrastructure is the capacity and collaborative mechanisms to specify, standardize, and share metadata, in a manner that balances autonomy and flexibility. For example, the use of distributed ontologies, that support real-time merging to discover needed datasets and their locations at a sufficiently granular level to efficiently take action (e.g., move, compute, analyze, or visualize). Online registries will be in place to enable researchers to access the contents of remote databases, formulate queries, transfer data, and combine disparate data management systems such as health, economics, etc. It then will become feasible to issue data requests that access many different data archives simultaneously, while rendering the underlying systems complexity transparent to the user. This model also can be extended to the sharing of applications, libraries, and software modules as follows:

- Support services for collaborative visualization, where users can make scientific understanding of the data via visualization tools in a collaborative environment.
- Support for distributed data discovery through ontologically-based metadata (i.e. Resource Description Framework and Web Ontology Language), including geo-spatial metadata
- Support for computation, analysis, and management of distributed data
- Support for application/module sharing, including virtualization strategies (e.g., diagnostic libraries).

Storage Management and Data Movement Issues and Challenges

The data volume expected from future simulations and observational devices is so large that it presents special challenges in the areas of storage management and large-scale data movement. While we foresee that storage technology will be able to store petabytes of data on disks, the mirroring of such data presents the following requirements:

1. Moving a petabyte of data, even at the rate of 1 GB/s (10 Gb/s), takes 10^6 seconds, or 277 hours (about 12 days). The entire end-to-end system must support that transfer rate and must be robust and able to recover from transient failures of the storage systems and networks involved.
2. All storage systems should interoperate seamlessly. Since storage systems are diverse (based on various standards and storage components), that goal requires a standard interface to allow seamless access of data from online storage (disks or solid-state drives) and offline storage (robotic tape systems). Such standards are starting to emerge in the High Energy Physics communities (called Storage Resource Management) and could help mitigate this problem.
3. When data are replicated, mechanisms should be in place to manage the lifetime of the replicated data, so over time the storage systems can reclaim the space the replicas occupy, while still preserving the "master" copies.
4. It is necessary to keep rich metadata of the data holdings content based on agreed upon terminology to search the data (referred to as "ontologies" or "controlled vocabularies"). It is also necessary to keep track of the locations of replicated data, and propagate version changes to the replicas. Furthermore, algorithms should be developed to allow users to get data from the sites that maximize transfer efficiency based on network traffic, storage system transfer capacity, and storage system load.

Data Analysis Center Issues and Challenges

Data replication allows data to be closer to users to reduce the latency of data movement. However, it is unreasonable to require that each user will download large volumes of data to perform the analysis, as is the practice now. Given the large volume of data, it may be more practical to perform the analysis tasks close to where the data are stored. This can be achieved by having "analysis centers" that are placed close to the data (Figure 1). The analysis center facilities (usually small computer clusters) should be designed to permit users to invoke their preferred analysis components and introduce their own analysis codes into such facilities. We believe that the model of "analysis centers" will simplify the tasks of climate analysts and avoid moving and replicating data unnecessarily. The analysis centers also should provide software components that compare model output in models have different grids, different resolution, and compare/combine data coming from different sites. Finally, the analysis centers should be designed to take advantage of multiple processors and multiple cores to run parallel analysis whenever possible. This requires an effort to parallelize the analysis tools used most by the climate community.

Data Environment and Workflow Management Issues and Challenges

Data analysis usually involves multiple steps. An analysis center would benefit from using workflow management systems to orchestrate these steps. Furthermore, the workflow process, and the details on the parameters used, what input data was used, and what was produced needs to be captured. Such information is referred to as "provenance." A similar need exists when running simulations, where a workflow and all the conditions for running the simulations need to be captured, so that simulations can be reproducible. While in practice, large-scale simulations do not need to be rerun, the provenance data ensure that all information about the runs are captured, preserved, and easily retrieved.

The workflow management system should allow some tasks to be executed on-demand. For example, in the above use case the user ran several climate models in real-time that generated multi-model ensembles for a desired region and period. It should be possible to have facilities that allow such on-demand runs. Here again, establishing new workflows or re-using existing ones should be possible.

Another capability that we believe will be most valuable to the climate community is the ability to support tagging data by community users. Such annotations have proved extremely useful as a mechanism to inform all users of observation of interesting patterns, anomalies, etc. Related to such services, is to allow users to see changes made to the data holdings and the history of such versions.

Analysis Issues and Challenges

Intelligence analysis, emergency response, disaster prevention, and border security concerns will play an important role in the shaping of advanced analysis and visualization tools. A massive increase in climate simulation capacity will demand major advances in our analysis and visualization tools to elevate scientists' productivity. Many of the existing tool-development efforts (e.g. CDAT, NCL, GrADS, Ferret, etc.) are grossly understaffed relative to future challenges; however, the potential exists for an effective interagency collaboration to address this problem. The current tools already are inadequate for our high-resolution runs and are going to become even more problematic in the near future. These tools need to advance from single-threaded desktop applications to more powerful ensemble analysis capabilities that can leverage the same systems that produced the data. Analysis capabilities must scale along with the simulation capabilities.

The scale of data that needs to be analyzed will require that many analysis functions be performed near the data. For example, extracting desired data subsets would require applying indexing technology at the data source. Similarly, summarization operations (such as monthly means) will have to be performed at the data source. In general, multi-step coordination of the analysis pipeline will have to be managed using automated workflow capabilities that can control distributed processing.

Climate analysis is based on climate models and their verification with observed datasets. One of the most difficult challenges is comparing large-scale simulations with observed data because of the differences in scale and layout of the data. Observed data can originate from various devices, such as monitoring stations and satellite images. The analysis tools that integrate and compare such data with simulated data will present unprecedented challenges as the scale and diversity of the data grows over time.

In the next few years, non-experts will increasingly be analyzing climate data. This will make the strong case for using other data formats such as Geographic Information System and Google Earth Keyhole Markup Language and conversion tools. This effort will require richer metadata standards to provide more comprehensive information.

Visualization Issues and Challenges

Scientific visualization, which is the science and art of transforming abstract data into readily comprehensible images, is an integral and indispensible part of the modern scientific process. It leverages the high-bandwidth human visual system along with humans' ability to assimilate visual information into comprehension and understanding.

Visualization research and application efforts tend to fall into one of three primary use modalities. The first is "exploration visualization" where the user has no idea what they are looking for. This type of use model is typically the most challenging since its success relies on interactive, “random-access” exploration of large and complex datasets. Another use model is “analytical visualization” where the user knows what they are looking for in data. Analytical visualization techniques often are those reduced into optimal practice after being established during exploratory visualization. Finally, ”presentation visualization” is where the user wishes to convey a specific concept to others.

It is clear that the coming evolution in climate science will result in an increased ability to generate, collect, and store larger and more complex datasets. The size and complexity of these datasets will result in several challenges. First, these data will be of unprecedented spatial and temporal resolution with spatial resolution far exceeding screen resolutions. Second, many of the new simulation datasets will likely exhibit advanced spatial mappings like adaptively refined meshes (space and time) and icosahedral/geodesic grids. Third, future visual data analysis software infrastructure will be deployed in diverse ways to support several use modalities ranging from standalone-desktop applications to highly parallel tools that are integrated as part of easy-to-use, web-interface brokered visual data analysis infrastructure.

The climate science and visualization communities have done an admirable job over the past few decades in producing visual data analysis software infrastructure. However, we are at a crucial juncture marked by the growth of data size and complexity, the proliferation of parallel computational platforms, and the need for a large, global community of climate scientists to share data and software to solve some of our generation's most challenging and urgent problems. Some of the challenges facing visual data analysis at this juncture are as follows:

- Existing tools and algorithms are not suitable for use on larger datasets. Tools from the 1970s and 1980s that have been the “workhorses” of the past few decades simply will not run on many current and virtually all future large, complex datasets.
- Most existing tools/algorithms, in serial form, lack the capacity or responsiveness to meet the future visual data understanding needs of the climate community. For example, most existing tools and algorithms are serial. In contrast, it is generally accepted that the path towards the target capability will require effective use of parallel computational platforms and parallel data input/output infrastructure.
- The explosion of data size and complexity will give rise to a new set of needs in visual data understanding. Existing techniques, which have proven effective for single-variable, coarse datasets, will not be adequate for future science needs. For example, existing visual data analysis algorithms are not effective for studying the relationships between the dozens or hundreds of runs produced in a single ensemble or for comparative analysis of dozens or hundreds of ensemble runs. Other examples include the need to discover and understand the relationships between variables in time-evolving data, comparing simulation and experiment/observed data, and comparing/understanding data that exist on different grid types and have different resolutions.
- Software engineering to enable deployment of new capabilities in support of diverse use modes. The target here is to have visual data analysis software infrastructure that can be deployed as a standalone application on a desktop or as part of workflow that runs on large parallel computational platforms.
- Computational platforms and infrastructure for visual data exploration/analysis post-processing need to be designed to accommodate data intensive computing needs, which include vast amounts of input/output bandwidth and memory. These systems need to be available for on-demand, real-time visual data exploration to support the “exploratory” form of visual data analysis, as well as to support on-line execution of “analytical” and “presentation” visualization.

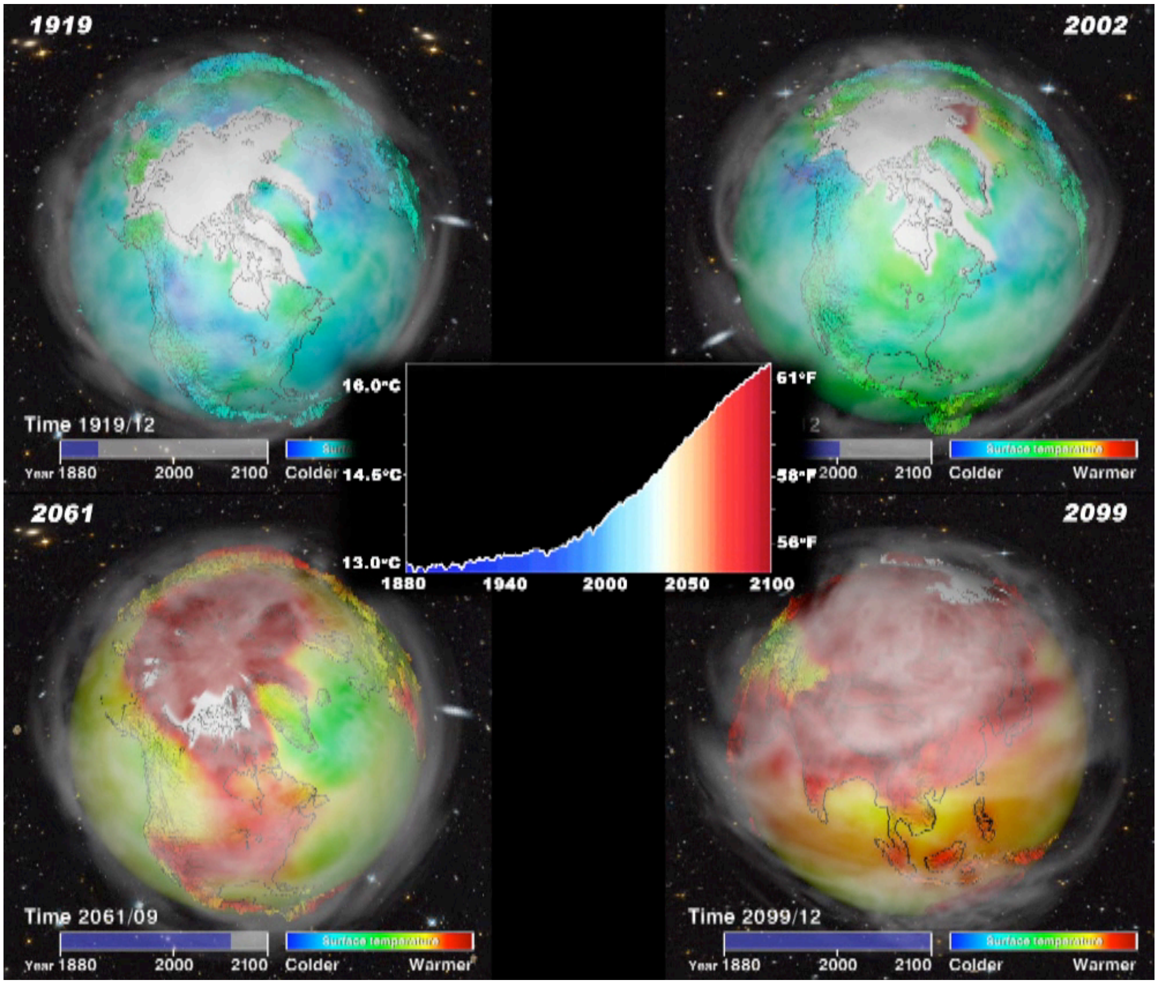

Figure 3. 220 years of science observation and computer simulated global warming.

Conclusions

Our objective is to envision the critical and unfolding “grand challenges” in data management, analysis, and visualization, and to advance the capacity for collaborative research communities to become high-performing stewards of all distributed, scientific workflow processes associated with extreme-scale datasets. This is a complex and exploratory process. Climate modeling as a community activity is becoming integral to an inter-dependent chain of operations in which global models force other models that predict societal impacts of climate, for example in agriculture, water resources, energy-usage, etc. The ability to quickly implement downscaled dependencies between such operational modeling chains must become routine. Extreme data volumes will tax any foreseeable network bandwidth. Therefore, it is imperative that robust multiple data reduction capabilities be developed and deployed at provider sites and readily provisioned to users via common standards.

In the short-term, much of the work we discuss continues to be neglected and needs immediate action if the community is to realize extreme-scale data management, analysis, and visualization for the future.

References

External Links

1. The ESG:
    a. "Earth System Grid Center For Enabling Technologies Documents" – http://esg-pcmdi.llnl.gov/documents
    b. "Earth System Grid Data Portal" – http://www.earthsystemgrid.org

2. Wide Area Network:
    a. "National Science and Technology Council - Federal Plan for Advanced Networking Research and Development" – http://www.er.doe.gov/ASCR/ProgramDocuments/Docs/ITFAN-FINAL.pdf
    b. "Workshop Report on Advanced Networking for Distributed Petascale Science: R&D Challenges and Opportunities" – http://www.er.doe.gov/ASCR/ProgramDocuments/Docs/NetworkResearchWorkshopReport08.pdf
    c. "Science Requirements for ESnet Networking Documents" – http://www.es.net/hypertext/requirements.html

3. Data Management:
    a. "DOE Office of Science Data-Management Workshop Report" – http://www.er.doe.gov/ASCR/ProgramDocuments/Docs/Final-report-v26.pdf

4. Visualization:
    a. "Visualization and Knowledge Discovery: Report from the DOE/ASCR Workshop on Visual Analysis and Data Exploration at Extreme Scale" – http://www.sci.utah.edu/vaw2007/index.html

5. Climate Analysis Tools:
    a. Climate Data Analysis Tools (CDAT) – http://www2-pcmdi.llnl.gov/cdat
    b. NCAR Command Language (NCL) – http://www.ncl.ucar.edu
    c. Grid Analysis and Display System (GrADS) – http://www.iges.org/grads
    d. Ferret – http://ferret.wrc.noaa.gov/Ferret

6. Software Development:
    a. "Software Development Tools for Petascale Computing Workshop" – http://www.er.doe.gov/ASCR/WorkshopConferences/sdtpc_workshop_report.pdf