CEDPS:
Center for Enabling Distributed Petascale Science
Management Plan

1. Introduction

The Center for Enabling Distributed Petascale Science (CEDPS) will produce technical innovations for (a) rapid and dependable *data placement* within a distributed high-performance environment, (b) the convenient construction of *scalable science services* that provide for the reliable and high-performance processing of computation and data analysis requests from many remote clients, and (c) *troubleshooting* of these and other related ultra-high-performance distributed activities, from the perspectives of both performance and functionality. The CEDPS program of work has been defined in consultation with leading DOE application groups, and will be deployed, applied, and evaluated in collaboration with major DOE projects in climate, high energy and nuclear physics, combustion, astrophysics, fusion, biology, and other sciences.

The project team brings together researchers and technologists from three DOE laboratories, Argonne National Laboratory, Fermi National Accelerator Laboratory, and Lawrence Berkeley National Laboratory, and two universities, University of Southern California, and University of Wisconsin. This team has an outstanding history of both successful research and technology development, as evidenced by such products as Globus, Condor, NetLogger, pyGlobus, and dCache, *and* successful transfer to the DOE science community, as evidenced by the key role that team members' software already plays in major DOE projects such as Earth System Grid and Open Science Grid.

Management of the center will be the overall responsibility of the Principal Investigator, Ian Foster and the Project Director, Jennifer Schopf, both at Argonne. The Project Director will be responsible for general project coordination and organization, tracking of cross-over issues, coordination of outreach activities, and moderating of priority discussions between application requests and technologists. Jointly with the PI, the project director will coordinate general project outreach and direction.

Each of the three focus areas has an Area Lead: Ann Chervenak, ISI, for Data, Kate Keahey, Argonne, for Scalable Services, and Brian Tierney, LBNL, for Troubleshooting. Each Area Lead will be responsible for milestones, vision, and coordination within their area.

Project tracking will take place on weekly teleconferences involving the Project Director, PI, and Area Leads, with other members of CEDPS welcome to attend as relevant. In addition, each area will have weekly or bi-weekly meetings as necessary.

Full project meetings will take place every 6 months. The first project meeting took place at ANL October 25-26, 2006, and the second will take place March 12-13 at USC. Annual reports will be submitted to DOE on approximately October 1 of each year, and interim reports will be submitted on approximately April 1 of each year.

We plan to create software releases approximately every 12 months, with improved service functionality All software will be released as Open Source with a FreeBSD-like license, and will be supported via a public email lists and a wiki at http://www.cedps.net.

2. Metrics for Success

In addition to the milestones listed for each area, the overall project will be evaluated by the following metrics of success.

1. How well is the outreach to application groups functioning as a mechanism for two-way exchange of information – both letting users know about the new functionality offered in CEDPS software and as a set of requirements drivers for development?

This will be evaluated by requesting feedback from each of the named applications groups for the midterm and annual reports, and an evaluation of contact made to new user groups. Any tutorials or outreach meetings will also be listed in the reports.

2. How well is CEDPS software and tools enabling SciDac end users to overcome the barriers of Petascale Science?

The most common way to measure this is the traditional count of papers produced by the science end users after deployment of the software, but we will also report directly from users on new functionality or approaches enabled by the CEDPS tools.

3. How well are the new data tools (MOPS, RDS, DPS) increasing the performance and usability of the SciDAC applications communities?

This will be answered by taking performance measurements throughout the projects work with applications groups, and evaluating improvement in terms of number of files managed, transfer times, number of replicas used, number of users on the system, and user feedback in terms of response time and usability. These values will be including in the reporting.

4. How many applications are successfully using the Scalable Service tools to adapt their existing applications into network services?

This will be measured through a count of deployments, and testimonials from user communities.

5. How much has the CEDPS troubleshooting infrastructure improved job throughput and file transfer reliability and performance for a set of targeted applications communities (starting with OSG and ESG)?

The key to this metric will be evaluating failure behavior both before and after CEPDS software is used. To date, we have only been able to reliably establish metrics once a deployment has occurred, so this value may show improvement with updates, as opposed to from a pre-CEDPS baseline.

6. How well is CEDPS enabling the scalability of applications on the DOE Leadership Class Facilities and other resources?

This will in part be measured by the number of deployments of CEDPS tools across SciDAC applications communities, reported yearly, but also by the number of interactions and deployments of software on selected environments.

3. Data

We will develop tools and techniques for reliable, high-performance, secure, and policy-driven placement of data within a distributed science environment. We will develop a set of end-to-end data distribution services that implement different data distribution and placement behaviors. To support these services, we will also construct a Managed Object Placement Service (MOPS)—a significant enhancement to today's GridFTP—that allows for management of the space, bandwidth, connections, and other resources needed to transfer data to and/or from a storage system.

3.1. Technology baselines for Data services

The data services that will be developed under CEDPS will extend previous work by CEDPS participants. The Managed Object Placement Service (MOPS), which will provide support for end-to-end management of data transfers, will extend and integrate existing work on the GridFTP data transport protocol, the NeST network storage device, the Stork data placement scheduler, the Storage Resource Manager, and the dCache system. The Reliable Distribution Service, which will implement mechanisms for reliable replication or placement of files, will extend and integrate functionality from the Reliable File Transfer Service, GridFTP, Stork and the Replica Location Service. The higher-level Data Placement Services, which will implement a set of strategies that determine which files should be placed at which locations, will extend the Data Replication Service and integrate with the Replica Location Service.

Globus GridFTP Family of Tools: GridFTP is a protocol defined in the Global Grid Forum (GGF) that provides secure, fast, efficient, and robust transport of data, particularly bulk data. Key features of the current implementation include the use of the eXtensible Input/Output (XIO) System, a Data Storage Interface (DSI), striping support, the server side processing and dynamic reconfiguration, robustness / restart improvements, and modularization to allow for easy extensibility. Documentation is at http://www.globus.org/toolkit/docs/4.0/data/gridftp/.

Reliable File Transfer Service (RFT): While globus-url-copy and GridFTP in general are a very powerful set of tools, additional functionality is needed for long transfers of many files. RFT provides "job scheduler"-like functionality for data movement that can recover from a wide variety of failures, including network outages, server failures, and client-side software or ardware failures. Documentation is at http://www.globus.org/toolkit/docs/4.0/data/rft/.

Replica Location Service (RLS): The Replica Location Service (RLS) is a distributed replica catalog that allows users to register the existence of replicas and discover them. Our distributed RLS consists of two components: a Local Replica Catalog that stores mappings from logical names of data items to the addresses of those data items on a storage system and a Replica Location Index that contains information about the mappings stored in one or more LRCs and answers queries about their contents. Documentation is at http://www.globus.org/toolkit/docs/4.0/data/rls/.

Data Replication Service (DRS): The Data Replication Service (DRS) provides higher-level data management functionality that gives users the capability of performing integrated data replication and catalog registration operations using a single service. DRS builds on lower-level Grid data services, including the Globus Reliable File Transfer (RFT) service, Replica Location Service (RLS), and GridFTP. The function of the DRS is to ensure that a specified set of files exist on a storage site. The operations of the DRS include *discovery*, identifying where desired data files exist on the Grid by querying the RLS; *transfer*, copying the desired data files to the local storage system efficiently using the RFT; and *registration*, adding location mappings to the RLS so that other sites may discover newly-created replicas. Documentation is at http://www.globus.org/toolkit/docs/4.0/techpreview/datarep/.

NeST: NeST is a software network storage appliance designed to provide guaranteed storage allocation. These allocation units, called *lots*, provide a guaranteed space for a period of time. Information is at http://www.cs.wisc.edu/condor/nest/.

Stork: Stork is a data placement scheduler, designed to optimize the scheduling of data movement. Stork allows for data placement jobs to be scheduled and managed, while providing for multiple transfer mechanisms and retries in the event of transient failures. With Stork now integrated into the Condor system, Stork jobs can be managed with Condor's workflow management software (DAGMan). Information is at http://www.cs.wisc.edu/condor/stork/.

Storage Resource Manager (SRM) : SRMs provide protocol negotiation, dynamic transfer URL allocation, advanced space and file reservation, and reliable replication mechanisms. The range of applications includes a permanent tape storage coupled with a distributed disk cache, SRM managed disk storage, or a distributed disk storage. Information is at http://computing.fnal.gov/ccf/projects/SRM/.

dCache: The dCache software system manages individual disk storage nodes and makes them appears as a single storage space with a single file system root. It has an SRM v1 and v2 interface, supports GridFTP and other transports for whole-file data movement, and includes a proprietary POSIX-like interface (dcap) for random access to file contents. Documentation is at http://www.dcache.org/.

3.2. Milestones

Year 1

- Work with application communities, including OSG, ESG, LIGO, others, to understand their requirements and use these to drive service development. (AFIW)
- Enable local resource management within GridFTP server: Prototype and implement basic capabilities within the GridFTP server for protecting the underlying system from resource exhaustion, in particular, to better manage network connections and memory usage. (A)
- Prototype and implement NeST capabilities for managing the underlying storage component of GridFTP transfers. An important study will be to minimize the storage management impact on the GridFTP transfer rate. (W)
- Combine enhanced GridFTP and NeST services into an early MOPS 0.5 release with basic functionality. (AW) This release will be made available for users to start utilizing a basic MOPS functionality as soon as possible. We will work with OSG on evaluating, testing, and incorporating MOPS into the VDT. (W)
- MOPS 1.0 release by year end that builds on the MOPS 0.5 implementation and adds capabilities for exposing information on internally managed resources. This resource information can then be used by higher level services to more efficiently manage transfers in services such as RFT and to perform resource matching and scheduling in tools such as Stork. We note that exposing resource information is only a first cut as some resource information may not be available via this mechanism. More sophisticated logic may be required to handle this problem at a higher level. MOPS 1.0 bundles will also be targeted for release in an upcoming OSG VDT release. (W)
- Research optimization for transfer of Lots of Small Files (LOSF)
- Investigate the integration of SRM and dCache with GridFTP service into the Managed Object Placement Service. Develop a plan for this work. (AF)
- Design and prototype implementation of a higher-level, policy-driven placement service that interfaces with the reliable distribution layer. (I)
- Design interface between higher-level placement services and the reliable distribution layer and implement a prototype. (I)

Year 2

- Continue working with application communities (OSG, ESG, others) to understand requirements and deploy and test CEDPS services. (AFIW)

- Research and basic implementation of dynamic MOPS resource registration (A)
- Interface design and prototype implementation for Reliable Distribution Service (AI)
- Research non-striped connection management using NeST.
- Work with targeted opportunities such as combustion applications (AL)
- Develop managed storage for non-striped MOPS. (AW)
- Research impact of managed transfers on storage in a feature-rich provisioning environment. (F)
- Continue working with OSG to test MOPS, and incorporate OSG feedback back into MOPS.
- Research methods for incorporating troubleshooting into MOPS. (A)
- Research into managed QoS networks.
- Deliver MOPS 2.0 release that includes additional optimizations such as LOSF optimizations in GridFTP and RFT, dynamic resource registration, non-striped connection management using NeST, and basic troubleshooting instrumentation. This release will also include an enhanced interface for resource matching. (AW)
- Deliver interface design and prototype implementation of Reliable Distribution Service. (AI)
- Continue development of policy-driven placement services, supporting additional policies (e.g., push/pull, maintaining *n* copies of files, etc.). Release placement service version 1.0. (I)
- Evaluate the integration of troubleshooting techniques into placement services. (I)

Year 3

- Implement LOSF research for advanced LOSF capabilities into MOPS and RFT. (A)
- Develop RLS local file name translation MOPS modules. (W)
- Experiment with use of data services for exchanging troubleshooting data.(AI)
- Prototype dynamic stripe width determination per transfer. (A)
- Research bandwidth management via XIO drivers and Lambda station. (AF)
- Develop distributed (striped) management of shared storage resources. (AFW)
- Update the MOPS release and the Reliable Transport Layers to incorporate new and advanced technologies.
- Continued development of policy-driven placement services, including adaptive placement based on feedback from reliable distribution layer and troubleshooting services. Release placement service version 2.0. (I)
- Initial investigation of dynamic deployment of data services. (AFIW)
- Continue interactions with application communities (OSG, ESG, others) to understand requirements and deploy and test CEDPS services. (AFIW)

Years 4-5

- Develop an updated plan for years 4 and 5 that takes into account current technologies and application requirements.
- Deliver production releases of data services: Placement Services, Reliable Transfer Services, MOPS (AFIW)
- Continue to deploy and support data services in OSG, ESG, other application environments (AFIW)
- Integration of troubleshooting into data services (AFILW)
- Integrate data services with the Application Hosting Service (AHS), experiment with dynamic deployment of data services. (AI)
- Demonstrate integrated data placement at scale and with provisioning and failures (AFIW)

3.3. Roles and Responsibilities within Data

- Ann Chervenak, ISI (25% on CEDPS), is the Data Area Lead. She is also responsible for the design and direction of work on the Data Placement Service, and primary point of contact between the Data area and ESG and LIGO.
- Rob Schuler, ISI (50% on CEDPS), is an architect and developer who will be responsible for the design and implementation of the Data Placement Service, as well as interfaces to Reliable Distribution Service.
- Dan Fraser, ANL (35% on CEDPS), is responsible for the design and direction of work on MOPS, and leads the GridFTP team. He is the primary point of contact between the Data area and OSG, and will be involved in future Data area outreach.
- John Bresnahan, ANL (40% on CEDPS), is a developer and data transport architect for GridFTP and Globus data management software. He will assist with the design, development and deployment of the MOPS, DPS, and RDS services.
- Mike Link, ANL (40% on CEDPS), is a developer and data transport expert for GridFTP and Globus data management software. He will assist with the design, development, and deployment of the MOPS, DPS and RDS services.
- Raj Kettimuthu, ANL (30% on CEDPS) is a developer and QoS expert for Globus data management software. He will assist with the design, development, and deployment of the MOPS, DPS and RDS services.
- Ravi Madduri, ANL (10% on CEDPS), is the lead developer for RFT, and will assist on the design, development, and deployment of the MOPS, DPS and RDS services.
- Miron Livny, Wisconsin (8% on CEDPS), leads the CEDPS effort at Wisconsin and is responsible for technical guidance and management.
- Nick LeRoy, Wisconsin (75% on CEDPS), is a developer with responsibility for enhancing NeST maintaining NeST and Stork.
- Fang Cao, Wisconsin (50% on CEDPS), will work on Stork development
- Gene Oleynik, FNAL (10% on CEDPS), leads the CEDPS effort at Fermi Lab.
- Mark Bowden, FNAL (20% on CEDPS), works on SRM and dCache development at Fermi Lab.

- Alexander Moibenko, FNAL (25% on CEDPS in year two), is the lead integrator of Lamda Station with dCache/SRM at FNAL.
- TBD, FNAL (50% on CEDPS), for additional development on SRM and dCache.

4. Scalable Services

Providing an application as a remotely accessible service greatly simplifies community access to it by enabling clients to use it without having to install the code themselves or provision resources for its execution. Therefore we are developing within CEDPS infrastructure that will make such application hosting possible, consisting of a suite of services called the Application Hosting Services (AHS).

The AHS developed in this project will consist of two major components. The front-end, developed using pyGridWare will generate Web Service interfaces turning existing applications into network services. On the backend, we will develop a range of services for flexibly provisioning resources to ensure on-demand service deployment in response to time varying requirements. Those services will be implemented using GRAM and the Workspace Service.

4.1. Technology baselines

We will rely on a number of technologies in developing the infrastructure for this task. Specifically, we will use pyGridWare to automatically generate Web service interfaces for applications, the Globus GRAM service to manage jobs, and the Globus Workspace Service to manage physical and virtual resources. Under CEDPS we will to extend and improve these tools and develop infrastructure to integrate them.

Python Grid Tools: PyGlobus and pyGridWare provide the ability to develop both pre-Web Service and WSRF based Grid clients and servers. Both toolkits support the easy integration of codes written in C/C++ and Fortran and support the rapid development of science services. The Python Command Line Service Tool (PyCLST) supports wrapping simple command line applications to expose them as Grid Services. PyCLST allows application providers to easily expose their command-line applications as Grid Services without having to understand the details of WSRF and WS-N. Documentation is at http://dsd.lbl.gov/gtg/projects/pyGlobus/.

GRAM: The Globus GRAM service was developed to facilitate processing requests for remote application execution and managing active jobs. GRAM became the protocol interaction between the administrators of each application and the administrators of the available resources. To provide this functionality, GRAM interfaces with local resource managers such as PBS or LSF. Documentation is at http://www.globus.org/toolkit/docs/4.0/execution/wsgram/.

Workspace Service: The Globus Workspace Service allows a Grid client to deploy and manage workspaces – execution environments that can be dynamically deployed in the Grid. The current implementation focuses primarily on the deployment and management of virtual machines implemented using the Xen hypervisor. The Workspace Service allows the clients to manage the resources assigned to each environment. Documentation is at http://www.globus.org/toolkit/docs/4.0/techpreview/wms/.

4.2. Milestones

Year 1

- Develop an architecture document for dynamic resource provisioning using GRAM; demonstrate a prototype and gather user requirements from ESG and others.

- Develop a proof-of-concept infrastructure for dynamic provisioning in response to time varying application requirements using the Workspace Service; apply the service to LBL STAR application and potentially others; deploy the infrastructure on at least two sites and gather requirements based on preliminary user feedback.
- Based on proof-of-concept experiences, develop an architecture and roadmap for dynamic provisioning using virtual machine technology; initiate its implementation.
- Initiate the development of “Workspace Service bookshelf”: library of workspace images for scientific applications.
- Release a new version of the Workspace service integrating features developed in response to proof-of-concept requirements.
- Identify “network service” requirements for Web service application wrapping infrastructure using pyGridWare for Fusion applications and others; build proof-of-concept and prototype applications.
- Develop pyGridWare-based portal generation services.
- Develop an initial version of pyGridWare instrumented with the NetLogger toolkit.
- Develop an instrumented version of the Python Command Line Service Tool.
- Develop adaptor classes that support the pyGlobus interfaces to GRAM and other services based on pyGridWare.
- Implement additional services in pyGridWare to support and information services, and delegating proxies to the container.
- Deploy new services on OSG and ESG.
- Work with Grid application developers to assist them with deployment of developed capabilities.

Year 2

- Develop a preliminary architecture document integrating the Web service application infrastructure with provisioning backends.
- Work with biology applications on creating science services using initial AHS (job management based as well as resource management based solutions wherever appropriate).
- Develop protocols for specifying targets for scalable services, including performance and resource provisioning targets; continue the implementation of workspace-based provisioning.
- Further work on “workspace bookshelf”: developing schemas for describing and identifying execution environments.
- Release the first version of services for on-demand provisioning of workspaces.
- Continued developed of pyGridWare to support new protocol versions.
- Develop a version of PyCLST that supports wrapping non=command line applications.
- Continue to deploy new services on OSG, ESG, and others.
- Continue to work with Grid application developers to instrument their applications

Year 3

- Integrate troubleshooting with AHS services and release a version of the Workspace Service reflecting it.
- Continue to deploy new services on OSG, ESG, and others.
- Continue to work with Grid application developers to instrument their applications

Years 4-5

- Integrate the Web service interfaces/application portals with the AHS infrastructure.
- Continue to deploy new services on OSG, ESG, and others.
- Continue to work with Grid application developers to instrument their applications

4.3. Roles and Responsibilities within Scalable Services

- Kate Keahey, ANL (30% on CEDPS), is the Scalable Services Area Lead. She is also responsible for the design and direction of scalable services efforts to provide provisioning for services, and the primary contact for STAR and GADU.
- Tim Freeman, ANL/UC (40% on CEDPS), is a software developer working on the Workspace Service.
- Carl Kesselman, ISI (10% on CEDPS), is the lead on GRAM-based approach to scalable services, and the primary scalable services contact for ESG. He is also the lead for CEDPS work at ISI.
- Keith Jackson, LBNL (50% on CEDPS), is the lead on the Web service application wrapping effort, and the primary contact for FUSION and DANCE. He is also the lead for CEDPS work at LBNL.
- David Konerding, LBNL (50% on CEDPS), is a software developer working on pyGridWare and PyCLST.

5. CEDPS Troubleshooting

Grid Users and administrators must be able to detect errors at run time, analyze baseline performance shifts, and receive warnings about failures. Therefore Grid troubleshooting is one of three areas being addressed by CEDPS. We will be developing tools to assist in the automated failure detection for Grid Services, the end-to-end failure detection and analysis of user jobs, including log file discovery and management, and automated performance degradation detection for both Grid Services and user job.

5.1. Technology baselines

The main technologies we are starting with for this task are the MDS4 Trigger Service and the NetLogger Toolkit. MDS4 Trigger Service provides the ability to generate email alerts based on resource monitoring information. NetLogger provides a set of tools to aid in the instrumentation of software, tools to collect the log files, and tools to analyze the logs. Each of these baseline technologies will be greatly improved by the CEDPS project.

Trigger Service: The Globus MDS4 Trigger service periodically queries the MDS4 Index Service to obtain resource information and compares that data against a set of conditions. When a condition is met, the Trigger service performs an action, such as emailing system administrators to notify them that a monitored component is down or that disk space on a server has reached a threshold. The current functionality will be expanded and generalized for use beyond the current

ESG deployment, starting with OSG. Documentation is at http://www.globus.org/toolkit/docs/4.0/info/trigger/.

NetLogger: NetLogger log file management tools are designed to handle the collection of NetLogger formatted log files. The tools are unique in that they are targeted at high-volume logs typical of application instrumentation. These will be generalized to handle all types of log files, and integrated with the MDS4 and the new log archive service. NetLogger also include some prototype anomaly detection tool to locate missing workflow events based on a predefined list of expected events. This will be expanded upon and integrated with MDS4. Documentation is at http://dsd.lbl.gov/NetLogger/.

We will also incorporate existing application-level and system-level instrumentation such as *syslog* and and *syslog-ng* into our framework

5.2. Milestones

Year 1

- Version 1.0 of the log file Collection Service: This service can be used to gather middleware and application log files from multiple locations to a single host.
- Instrumentation "best practices guide": This will be advice to middleware and application developers on what to log for maximal troubleshooting abilities. This includes what data should be included in the log files, and what format the log file should be. It will also include information on how to best generate and log "activity IDs" to aid in tracing the workflow of a Grid job.
- Version 1.0 of the log Archive Service: The Archive Service is necessary to be able to track baseline performance, monitor hardware and software configuration changes, and perform anomaly detection.
- Instrument Globus and Condor based on the new "best practices guide".
- Develop additional MDS4 Trigger Service Information Providers.
- Develop additional MDS4 Trigger Service action scripts (currently only send email).
- Deploy new services on OSG and ESG.
- Outreach to Grid application developers to instrument their applications.

Year 2

- Add authentication and authorization capability to the log Collection Service and log Archive Service.
- Tools to filter and feed log data from the Collection Service to the Archive Service.
- Use the Archive Service to establish performance baselines, and trigger events if performance deviates too much from the baseline.
- Develop MDS4 Trigger Service action scripts to securely restart failed services.
- MDS4 Triggers for missing log events (based on NetLogger anomaly detection tool).
- Integration of Log Collection Service with MDS4 to provide a log file location service.
- Continue to deploy new services on OSG, ESG, and others.
- Continue outreach to Grid application developers to instrument their applications

Year 3

- Release new troubleshooting Analysis Tools targeted to specific OSG, ESG, and specific application needs.
- Release new Trigger Service functionality.
- Release Data Mining Toolkit for building new analysis tools from monitoring data that follows the instrumentation "best practices" guide.
- Continue to deploy new services on OSG, ESG, and others.

Outreach to Grid application developers to instrument their applications

Years 4-5

- Data mining tools to look for sources of problems, integrated with the MDS4 Trigger Service.
- Integration of Trigger Service with application portals and workflow engines.
- Continue to deploy new services on OSG, ESG, and others.
- Continue to work with Grid application developers to instrument their applications

5.3. Roles and Responsibilities within troubleshooting

- Brian Tierney, LBNL (40% time on CEDPS), is the Troubleshooting Area Lead. He is the lead interface with OSG on requirements gathering and deployment efforts.
- Dan Gunter, LBNL (35% time on CEDPS), is a researcher and software developer with lead responsibility for the log collection service and the analysis tools.
- Jennifer Schopf, ANL (15% on CEDPS), is responsible for requirements gathering and interfacing with various applications communities. She will also interface to the MDS4 development group for shared effort with that work. She is also the lead for CEDPS work at ANL.
- TDB, ANL (20% on CEDPS), will be responsible for development and deployment of MDS4 related services, and augmenting Globus services for additional logging and troubleshooting work.
- Stu Martin, ANL (unfunded on CEDPS) is responsible for working with the GRAM auditing logs, and exposing data needed for troubleshooting. This is overlap work with TeraGrid and OSG.
- Laura Perlman, ISI (20% on CEDPS) is the lead interface with ESG on requirements gathering and deployment efforts. She is also a member of the MDS team.

6. Applications

CEDPS participants have a long history of collaboration with a wide variety of SciDAC and NSF applications. We have identified OSG and ESG as good drivers for our early work, due to their aggressive technology goals and complementary technology development/adoption schedules. We are also working with the HEP STAR application on dynamic service provisioning. We have begun discussions with a second wave of projects, including LIGO, FACETS, DANCE, Combustion, and GADU. We will expand this set over time via a sustained outreach program to the DOE community.

OSG's assertive role in attracting additional applications to their own community and the strong personnel ties between their project and ours makes it an attractive project for us to work with. Initial work with the facilities group, led by Livny, will be focused on log management from the CEDPS troubleshooting area. We are coordinating with them on a log management plan and best practice document, and they are the planned alpha testers for this work. In addition, we will work with OSG to integrate our end-to-end data movement tools and scalable service tools into its VDT software distribution, and then to deploy these tools across OSG sites. We will simultaneously engage with various OSG communities, in particular ATLAS and CMS, to incorporate those tools into their software suites.

We will engage with the climate community via ESG, with which we have already worked closely for five years. In the short term, we will extend their troubleshooting and error detection framework as a joint collaboration. To aid in the publication of climate model data produced at ORNL, LBNL, LANL, and NCAR, we will closely discuss development and deployment of MOPS so that the initial prototypes meet ESG's requirements. We hope that their testbed will be one of the first to deploy the higher-level data tools. In the longer term, we will use our scalable science service tools to create scalable server-side analysis services for this data and IPCC AR5 data. We have worked with the ESG-CET proposal to define mutually agreeable consistent deliverables and milestones that address these requirements.

To focus requirements on the dynamic provisioning side in the scalable services effort we initiated collaboration with the High Energy Physics STAR experiment. The STAR application is currently hard to install and port, thus making dynamic cycle provisioning for this application difficult if not impossible. Our initial efforts are focused on using virtual machines to install the application and deploy it on-demand on dynamically provisioned resources by using the Workspace Service. We expect this effort to provide a proof-of-concept of this mode of resource provisioning, and will use the prototype to focus future design and development work. We are fortunate to partner in this effort with the rPath company currently collaborating with LBNL on producing images for scientific application under a DOE SBiR contract. We expect to leverage these and other efforts in providing dynamic deployment capabilities to a wider set of applications.

Growing from these continuing collaborations, we have planned a strong outreach effort for the first 3 months of the project in order to extend our application interactions. This set, and the CEDPS member responsible for contact, includes:

- LIGO (OSG) – Ann Chervenak, data focus
- Fusion- FACETS – Keith Jackson, scalable services focus
- DANCE – Keith Jackson, scalable services focus
- Combustion – Brian Tierney, troubleshooting focus
- GADU – Kate Keahey, scalable services and data focus

In addition, we are attending meetings in the areas of eScience and distributed petascale computing, and hope to use these meetings to attract additional collaborations. This effort will be tracked at the weekly management calls, and CEDPS members will regularly attend other project meetings to encourage adoption and collaboration.

We also note that the popular software produced by CEDPS participants provides them with extremely strong connections with large and vibrant user communities that span DOE and other laboratories, US universities, and international science projects. These communities provide exceptional channels for requirements, code testing, user feedback, and even code contributions.

Our strong record of engagement in standards processes provides another avenue for engagement with the community at large.

7. Work with other Centers

In addition to these application-oriented collaborations, we will explore interactions with other SciDAC Centers and Institutes. In the first year, we will focus our interactions on three centers: Scientific Data Management Center for Enabling Technologies (SDM), PI Arie Shoshani, Lawrence Berkeley National Laboratory; Scaling the Earth System Grid to Petascale Data Center for Enabling Technologies (ESG), PI Dean Williams, Lawrence Livermore National Laboratory; and Sustaining and Extending the Open Science Grid: Science Innovation on a PetaScale Nationwide Facility (OSG), PI Miron Livny, University of Wisconsin.

Schopf and Foster will also be contacting and setting up meetings with other SciDAC application groups, and a visit to ORNL to increase collaborative ties to that group and the users of the Leadership facility at ORNL is planned for December 2006.