

The Grid Portal Development Kit

J. Novotny*,¹

¹ *Lawrence Berkeley National Laboratory, Berkeley, CA. 94704 U.S.A*



SUMMARY

Computational science portals are emerging as useful and necessary interfaces for performing operations on the Grid. The Grid Portal Development Kit (GPDK) facilitates the development of Grid portals and provides several key reusable components for accessing various Grid services. A Grid Portal provides a customizable interface allowing scientists to perform a variety of Grid operations including remote program submission, file staging, and querying of information services from a single, secure gateway.

The Grid Portal Development Kit leverages off existing Globus/Grid middleware infrastructure as well as commodity web technology including Java Server Pages and servlets. The design and architecture of GPDK is presented as well as a discussion on the portal building capabilities of GPDK allowing application developers to build customized portals more effectively by reusing common core services provided by GPDK.

KEY WORDS: Grid Computing; Computing Environments; Grid Portals; Science Portals; Globus

1. Introduction

Computational Grids [17] have emerged as a distributed computing infrastructure for providing pervasive, ubiquitous access to a diverse set of resources ranging from high-performance computers (HPC), tertiary storage systems, large-scale visualization systems, expensive and unique instruments including telescopes and accelerators. One of the primary motivations for building Grids is to enable large-scale scientific research projects to better utilize distributed, heterogeneous resources to solve a particular problem or set of problems. However, Grid infrastructure only provides a common set of services and capabilities that are deployed across resources and it is the responsibility of the application scientist to devise methods and approaches for accessing Grid services.

*Correspondence to: Jason Novotny (jdnovotny@lbl.gov), Lawrence Berkeley National Laboratory, 1 Cyclotron Rd., Berkeley, CA, USA 94704



Unfortunately, it still remains a daunting task for an application scientist to easily “plug into” the computational Grid. While command line tools exist for performing atomic Grid operations, a truly usable interface requires the development of a customized problem solving environment (PSE). Traditionally, specialized PSE’s were developed in the form of higher-level client side tools that encapsulate a variety of distributed Grid operations such as transferring data, executing simulations and post-processing or visualization of data across heterogeneous resources. A primary barrier in the widespread acceptance of monolithic client side tools is the deployment and configuration of specialized software. Scientists and researchers are often required to download and install specialized software libraries and packages. Although client tools are capable of providing the most direct and specialized access to Grid enabled resources, we consider the web browser itself to be a widely available and generic problem solving environment when used in conjunction with a Grid portal. A Grid portal is defined to be a web based application server enhanced with the necessary software to communicate to Grid services and resources. A Grid portal provides application scientists a customized view of software and hardware resources from a web browser.

Furthermore, Grid Portals can be subdivided into application-specific and user-specific portal categories. An application specific portal provides a specialized subset of Grid operations within a specific application domain. Examples of application specific portals include the Astrophysics Simulation Collaboratory [4] and the Diesel Combustion Collaboratory [34]. User portals generally provide site specific services for a particular community or research center. The HotPage user portal [36], the Gateway project [30], and UNICORE [35] are all examples of user portals that allow researchers to seamlessly exploit Grid services via a browser-based view of a well defined set of Grid resources.

The Grid Portal Development Kit [23] seeks to provide generic user and application portal capabilities and was designed with the following criteria:

- The core of GPDK should reside in a set of generic, reusable, common components to access those Grid services that are supported by the Globus toolkit [20] including the Grid Security Infrastructure (GSI) [21]. As Globus [18] becomes a de facto standard for Grid middleware and gains support within the Global Grid Forum [19], the GPDK shall maintain Globus compatibility through the use of the Java Commodity Grid (CoG) kit [32]. An enumeration and description of the Grid services is provided in the next section.
- Provide a customizable user profile that contains user specific information such as past jobs submitted, resource and application information, and any other information that is of interest to a particular user. GPDK User profiles are intended to be extensible allowing for the easy creation of application portal specific profiles as well as serializable such that users’ profiles are persistent even if the application server is shutdown or crashes.
- Provide a complete development environment for building customized application specific portals that can take advantage of the core set of GPDK Grid service components. The true usefulness of the Grid Portal Development Kit is in the rapid development and deployment of specialized application or user portals intended to provide a base set of Grid operations for a particular scientific community. The GPDK shall provide both an extensible library and a template portal that can be easily extended to provide specialized capabilities.

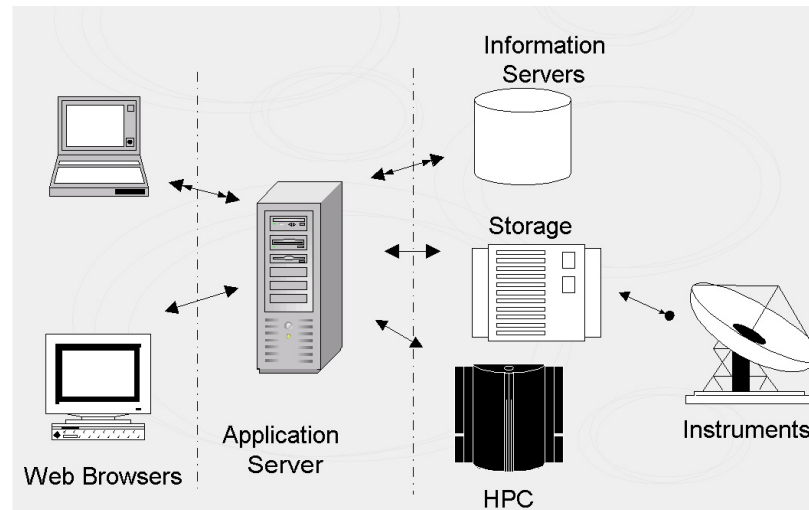


Figure 1. Standard 3-tier web architecture

- The GPDK should leverage commodity and open source software technologies to the highest degree possible. Technologies such as Java beans and servlets and widespread protocols such as HTTP and LDAP provide interoperability with many existing internet applications and services. Software libraries used by the GPDK should be freely available and ideally provide open source implementations for both extensibility and for the widespread acceptance and adoption within the research community.

The following sections explain the design and architecture of the Grid Development Kit with an emphasis on implementation and the technologies used. The advanced portal development capabilities of the Grid Portal Development Kit and future directions will also be discussed.

2. Overview of the Grid Portal Development Kit

The Grid Portal Development Kit is based on the standard 3-tier architecture adopted by most web application servers as shown in Figure 1. Tiers represent physical and administrative boundaries between the end user and the web application server. The client tier is represented as tier 1 and consists of the end-user's workstation running a web browser. The only requirements placed upon the client tier is a secure (SSL-capable) web browser that supports DHTML/Javascript for improved interactivity, and cookies to allow session data to be transferred between the client and the web application server.



The second tier is the web application server and is responsible for handling HTTP requests from the client browser. The application server is necessarily multi-threaded and must be able to support multiple and simultaneous connections from one or more client browsers. The Grid Portal Development Kit augments the application server with Grid enabling software and provides multi-user access to Grid resources. All other resources accessed by the portal including any databases used for storing user profiles, online credential repositories or additional resources forms the third tier, known as the back-end. Back-end resources are generally under separate administrative control from the web application server and subject to different policies and use conditions. The GPDK has been specially tailored to provide access to Grid resources as the back-end resources. It is generally assumed that Grid resources understand a subset of defined Grid and Internet protocols.

3. Grid Portal Architecture

The Grid Portal Development Kit provides Grid enabling middleware for the middle-tier and aids in providing a Grid enabled application server. The GPDK is part of a complex vertical software stack as shown in Fig. 2. At the top of the stack is a secure high-performance web server capable of handling multiple and simultaneous HTTPS requests. Beneath the web server is an application server that provides generic object invocation capabilities and offers support for session management. The deployed GPDK template portal creates a web application that is managed by the application server and provides the necessary components for accessing Grid services.

The Grid Portal Development Kit uses the Model-View-Controller (MVC) design pattern[29] to separate control and presentation from the application logic required for invoking Grid services. The GPDK is composed of three core components that map to the MVC paradigm. The Portal Engine (PE), provides the control and central organization of the GPDK portal in the form of a Java servlet that forwards control to the Action Page Objects (APO) and the View Pages (VP). The Action Page Objects form the “model” and provide encapsulated objects for performing various portal operations. The View Pages are executed after the Action Page Objects and provide a user and application specific display (HTML) that is transmitted to the client’s browser.

The Grid service beans form the foundation of the GPDK and are used directly by the Portal Engine, Action Page Objects and View Pages. The Grid service beans are reusable Java components that use lower-level Grid enabling middleware libraries to access Grid services. Each Grid service bean encapsulates some aspect of Grid technology including security, data transfer, access to information services, and resource management. Commodity technologies are used at the lowest level to access Grid resources. The Java CoG Toolkit, as well as other commodity software APIs from Sun and Netscape, provide the necessary implementations of Grid services to communicate a subset of Grid protocols used by the GPDK service beans.

The modular and flexible design of the GPDK core services led to the adoption of a servlet container for handling more complex requests versus the traditional approach of invoking individual CGI scripts for performing portal operations. In brief, a servlet is a Java class that implements methods for handling HTTP protocol requests in the form of GET and POST.

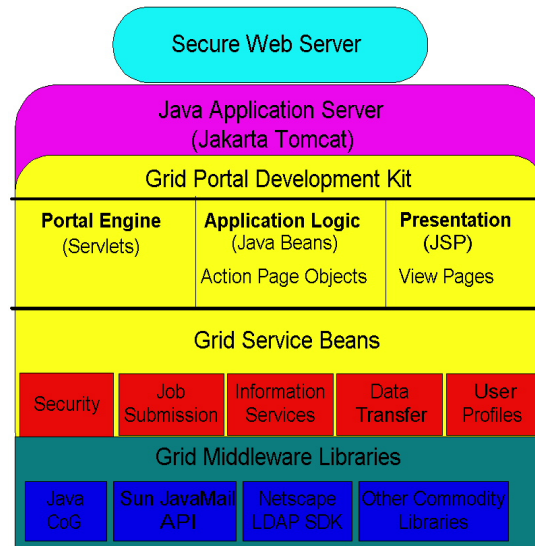


Figure 2. GPDK architecture and vertical stack of services and libraries

Based on the request, the GPDK servlet can be used as a controller to forward requests to either another servlet or a Java Server Page (JSP). Java Server Pages provides a scripting language using Java within an HTML page that allows for the instantiation of Java objects, also known as beans. The result is the dynamic display of data created by a Java Server Page that is compiled into HTML.

Figure 3 shows the sequence of events associated with performing a particular portal action. Upon start-up, the GPDK Servlet (GS) performs several key initialization steps including the instantiation of a Portal Engine (PE) used to initialize and destroy resources that are used during the operation of the portal. The PE performs general portal functions including logging, job monitoring, and the initialization of the portal informational database used for maintaining hardware and software information. The Portal Engine is also responsible for the authorizing users and managing users' credentials used to securely access Grid services. When a client sends an HTTP/HTTPS request to the application server, the GS is responsible for invoking an appropriate Action Page (AP) based on the "action value" received as part of the HTTP header information. The Page Lookup Table is a plaintext configuration file that contains mappings of "action values" to the appropriate Action Page Objects and View Pages. An AP is responsible for performing the logic of a particular portal operation and uses the GPDK service beans to execute the required operations. Finally, the GS forwards control to a View Page, a Java Server Page, after the AP is executed. The view page formats the results

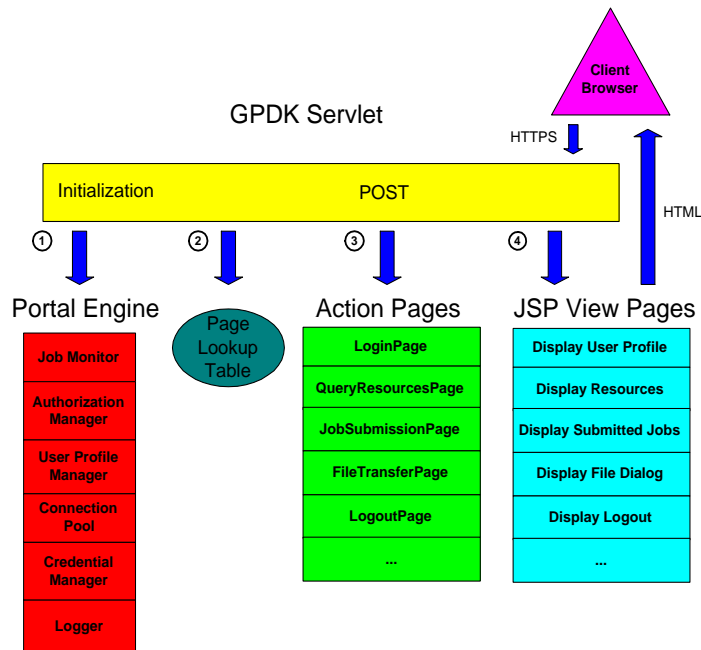


Figure 3. The GPDK Event Cycle

of an AP into a layout that is compiled dynamically into HTML and displayed in a client's browser.

4. GPDK Implementation

While a web server is unnecessary for the development of project specific portals using the GPDK, a secure web server is needed for the production deployment of a GPDK based portal. A production web server should offer maximum flexibility including the configuration of the number of supported clients, network optimization parameters, as well as support for 56 or 128-bit key based SSL authentication and support for a Java application server. The GPDK has been successfully deployed using the Apache [6] web server, a free, open source web server that provides a high-level of scalability and SSL support using the modSSL [7] package.

As mentioned previously, the GPDK relies on commodity Java technologies including Java beans, servlets and Java Server Pages for its general framework. The Grid Portal Development Kit was developed under the open source Tomcat application server available from the Jakarta Apache project [8]. The Tomcat [9] Application server was chosen as it is freely and widely



available and implements the latest JSP and Servlet specifications from Sun and is included as part of the Java Enterprise Edition (J2EE) production application server.

The Java Commodity Grid (CoG) toolkit provides most of the functionality required to implement the various Grid services that have been encapsulated by the core GPDK service beans. The Java CoG Kit was developed to provide compliance with Globus Grid services and compatibility with the C reference implementation of Globus. The Java CoG toolkit provides the following Grid protocols and services that are used by the GPDK:

- The Grid Security Infrastructure (GSI) provides a secure communication protocol that uses X.509 certificates and SSL to perform mutual authentication. The Java CoG Toolkit implements GSI using the IAIK [11] security libraries. Although the IAIK libraries are proprietary and not open source, they remain free for research and academic use. Implementation of GSI using other Java security libraries such as Sun's Java Secure Sockets Extensions [12] (JSSE) is being investigated.
- The Globus Resource and Management (GRAM) [16] protocol is used to submit jobs to a Globus gatekeeper, a standard authentication and job spawning service provided by Globus enabled resources.
- The Grid FTP [13] protocol provides an optimized data transfer library and a specialized set of FTP commands for performing data channel authentication, third-party file transfers, and partial file transfers.
- The Myproxy [33] service provides an online credential repository for securely storing users' delegated credentials. The Java CoG provides a client API for communicating to a Myproxy certificate repository for retrieval of a user's security credential.

One of the powerful features of Java beans in the context of web application servers and the GPDK service beans is *bean scope*. Bean scope refers to the level of persistence offered by a bean within the servlet container. For instance, beans may have session, application, or request scope. Session scope implies that the bean persists for the duration of a user's session, typically determined by the servlet container. For instance user profiles are represented as session beans and persist until a user decides to log out of the portal or their session times out as determined by the servlet container. Session scoped beans rely on the use of cookies used by most web browsers to retain state information on a particular client connection overcoming the inherent lack of state in the HTTP protocol. A bean with application scope persists for the complete duration of the servlet container and provides a persistent object used to store application specific static data that can be referenced by any Java Server Page on behalf of any user. The addition of collaborative capabilities such as a whiteboard or chat room, for instance, requires that messages be maintained with application scope, so logged in clients can see others' messages. A Bean with request scope persists only for the duration of a client HTTP request and is destroyed after the JSP page is processed into an HTML response for the client.

The GPDK has been developed and tested under Windows, Linux and Solaris platforms using the various JDK's provided by Sun in conjunction with the Apache web server available on both Windows and Unix platforms.



5. GPDK Services

The usefulness of the GPDK as a portal development framework rests on the currently supported set of common Grid operations that are required for a typical scientific collaboration. A scientific collaboration may involve one or more of the following capabilities:

- Submission, cancellation and monitoring of specialized programs (serial and/or parallel) to a variety of compute resources including those requiring batch (non-interactive) submission.
- The ability to store and retrieve data accumulated either from experiment or simulation to a variety of storage resources.
- Use resource discovery mechanisms to enable the discovery of hardware and software resources that are available to a particular scientific collaboration.
- The ability to perform the above operations securely by allowing scientists to authenticate to remote resources as required by the remote site administrators.
- Application specific profile information including user preferences, submitted jobs, files transferred and other information that scientists may wish to archive along with results obtained from computational simulations or laboratory experiments.

Within the GPDK framework, the above requirements have been encapsulated into one or more GPDK service beans. As discussed in the following sections, the GPDK service beans are organized according to Grid services in the areas of security, job submission, file transfer and information services. The deployed GPDK demo portal highlights the capabilities of the supported Grid services through template web and JSP pages.

5.1. Security

The security working group of Grid Forum has been actively promoting the Grid Security Infrastructure (GSI) [28] as the current best practice for securely accessing Grid services. GSI is based upon public key infrastructure (PKI) and requires users' to possess a private key and an X.509 certificate used to authenticate to Grid resources and services. A key feature of GSI is the ability to perform delegation, the creation of a temporary private key and certificate pair known as a proxy that is used to authenticate to Grid resources on a users behalf. The GSI has been implemented over the Secure Sockets Layer (SSL) and is incorporated in the Globus and Java CoG toolkit. One of the key difficulties in developing a portal to access Grid services is providing a mechanism for users to delegate their credentials to the portal since current web browsers and servers do not support the concept of delegation. Past solutions have involved the storage of users' long-lived keys and certificates on the portal. A user would then provide their long-term pass phrase to the portal, which creates a valid proxy that can be used on the user's behalf. The danger in this approach, however, is the risk of the web server being broken into and having possibly many users' long term private keys compromised. For this reason, the Myproxy service [33] was developed to provide an online certificate repository where users can delegate temporary credentials that can be retrieved securely by the user from the portal. Briefly, a user delegates a credential to the Myproxy server with a chosen lifetime and user name and pass phrase. The user would enter the same user name and pass phrase from their

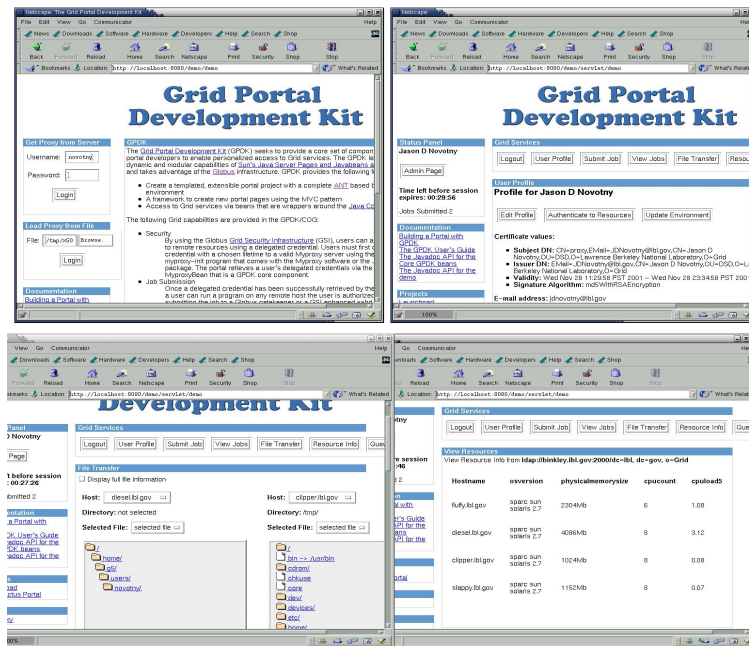


Figure 4. GPK demo pages clockwise from top left (a) login page (b) user profile page (c) resources page (d) file transfer page

browser over an HTTPS connection and the portal would retrieve a newly delegated credential valid for a chosen amount of time. Currently, GPK doesn't enforce any maximum lifetime for the credential delegated to the portal, but when a user logs off, the proxy is destroyed reducing any potential security risk of the delegated credential being compromised on the portal. The portal retrieves credentials from the Myproxy Server using the GPK security component, the MyproxyBean. The MyproxyBean component is actually a wrapper around the CoG toolkit client API to the Myproxy server. For users that have their delegated credential local to their workstation, the GPK template portal allows them to upload the proxy to the portal directly using standard file upload mechanisms over HTTPS. In addition, the JMyproxy package [24] provides a Java GUI that can create a proxy locally and delegate a credential to a Myproxy server. The initial login page that displays the Myproxy interface to the demo portal is shown in Figure 4(a). In the current implementation, all users that can either supply a delegated credential or retrieve one from the Myproxy server are authorized portal users. However, if the portal administrator wished to further restrict access, an encrypted password file on the portal or a secure back-end database could also be used to determine authorization information.



5.2. Job Submission

The GPKD Job Submission beans provide two different secure mechanisms for executing programs on remote resources. A GSI enhanced version of the Secure Shell (SSH) [21] software enables interactive jobs to be submitted to Grid resources supporting the GSI enabled SSH daemon. For all other job submissions, including batch job submissions, the Globus GRAM protocol is used and jobs are submitted to Globus gatekeepers deployed on Grid resources. Briefly, the GRAM protocol enables resource submission to a variety of resource scheduling systems using the Resource Specification Language (RSL) [16], allowing various execution parameters to be specified e.g. number of processors, arguments, wall clock or CPU time.

The primary GPKD components used to submit jobs are the `JobBean`, the `JobSubmissionBean` and the `JobInfoBean`. The `JobBean` provides a description of the job to be submitted, and encapsulates RSL by including methods for setting and returning values for the executable, additional arguments passed to the executable, number of processors for parallel jobs, batch queue if submitting a batch mode and more. The `JobSubmissionBean` is actually an abstract class that is sub-classed by the `GramSubmissionBean` in the case of submitting a job to a Globus gatekeeper or a `GSISSHSubmissionBean` using the GSI enhanced SSH client [21]. The `GramSubmissionBean` capabilities are provided once again by the Java CoG library.

Once a job has been successfully submitted, a `JobInfoBean` is created containing a time stamp of when the job was submitted and other useful information about the job, including a GRAM URL that can be used to query on the status of the job.

Job monitoring of submitted jobs is provided through the `JobMonitorBean`, a component initialized at start-up by the Portal Engine. The `JobMonitorBean` periodically queries the GRAM URL's on behalf of a user to keep track of job status based on the GRAM job status codes e.g. active, running, or failed. Because The `JobMonitorBean` has application scope, it can save job status to a user's profile even if the user has logged out.

5.3. File Transfer

The GPKD file transfer beans encapsulate the GridFTP [13] API implemented as part of the CoG toolkit and provide file transfer capabilities, including third-party file transfer between GSI enabled FTP servers, as well as file browsing capabilities. The `FileTransferBean` provides a generic file transfer API that is extended by the `GSIFTPTransferBean` and the `GSI SCPTransferBean`, an encapsulation of file transfer via the GSI enhanced scp command tool. The `GSIFTPServiceBean` provides a session scoped bean that manages multiple FTP connections to GSI enabled FTP servers. The `GSIFTPServiceBean` allows users to browse multiple GSI FTP servers simultaneously and a separate thread monitors server time-outs. The `GSIFTPViewBean` is an example view bean used by a JSP to display the results of browsing a remote GSI FTP server. Figure ? shows the demo file browsing and transferring page.



5.4. Information Services

The Grid Forum Information Services working group has proposed the Grid Information Services (GIS) architecture for deploying information services on the Grid and supported the Lightweight Directory Access Protocol (LDAP) as the communication protocol used to query information services. Information services on the Grid are useful for obtaining both static and dynamic information on software and hardware resources. The Globus toolkit provides an implementation of a Grid Information Service, known as the Metacomputing Directory Service using OpenLDAP, an open source LDAP server. Although, the Java CoG toolkit provides support for LDAP using the Java Naming and Directory Interface (JNDI), GPDK uses the open source Netscape/Mozilla Directory SDK [1] as it proved easier to use in practice and also provides support for developing a connection pool for maintaining multiple connections to several Grid Information service providers, thus eliminating the need for clients to re-connect during each query. However, this model will need to be re-evaluated with the widespread deployment of the MDS-2 architecture that includes GSI enhancements making it necessary for clients to re-authenticate to the MDS for each query. GPDK provides an `MDSQueryBean` and `MDSResultsBean` for querying and formatting results obtained from the MDS. Currently GDK supports querying the MDS for hardware information such as CPU type, number of processors and other details as well as CPU load and queue information that can be used by the user to make more effective job scheduling decisions.

5.5. GPDK User Profiles

User profiles are an integral part of the portal architecture as they enable the customization of a particular set of resources (hardware and software) by portal users. In addition, user profiles create a “value-added” feature of using the portal to perform common Grid operations since a user history is maintained in a user profile allowing users to keep track of past jobs submitted and results obtained. A user profile is a persistent session bean that is either created the first time a user logs into the portal or is loaded (de-serialized) from a file or database if one exists already. User profiles are serialized when a user logs out of the portal or the profile has been modified. Application portals built on top of the GPDK can subclass the GPDK user profile bean to provide additional application specific information for a particular class of users.

6. GPDK as a Portal Development Environment

Computational science collaborations can benefit significantly by providing web access to software, data and expertise. The Grid Portal Development Kit is designed to be a development environment that enables rapid development of application specific portals by leveraging off core GPDK service beans and the MVC architectural model incorporated by the GPDK. The core set of services remain generic enough to be useful to any community interested in staging data, executing codes, and providing a customizable, commonly accessible, collaborative environment.

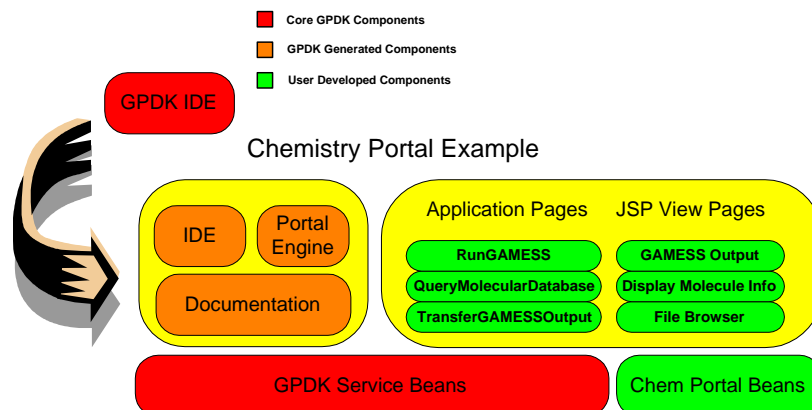


Figure 5. Creating a New Portal

The GPKD is packaged with a demo portal to showcase the functionality of the service beans and provide several common portal operations. During installation, the Grid Portal Development Kit deploys the library of core service beans as well as a central servlet and a collection of fully functional demo template web pages to the Java application server. The template web pages include HTML and Java Server Pages intended to demonstrate the GPKD service beans. The template source code contains the central servlet class used by the application specific portal as well as sub-classes for project specific user profiles and a project specific Portal Engine (PE) used to initialize and shutdown any resources required by the new project portal. The goal of the template portal is not to provide a fully polished interface ready for production deployment, but rather a demonstration of GPKD capabilities that can be easily extended and specialized for providing an application group with a common subset of Grid capabilities. Figure 5 shows the generalized GPKD architecture organized into core GPKD components, GPKD generated components, and application specific user developed components.

The GPKD template source code also contains all the Action Page Objects necessary for performing the various portal operations that demonstrate core GPKD services and beans. The APs are composed of a `LoginPage`, `UpdateProfilePage` and `LogoutPage` that demonstrates the retrieval of credentials from the Myproxy server and the loading, editing and saving of user profiles. The `JobSubmissionPage` and `FileTransferPage` demonstrate the GPKD service beans in the areas of job submission and file transfer. The template Java Server Pages are displayed upon the successful completion of the Page object operations. If an error occurs during the execution of an AP, an exception is thrown with a relevant error message that is displayed to the client by a specialized error JSP.



The Grid Portal Development Kit relies on ANT [2] for cross-platform project compilation and deployment. ANT provides an XML based Java build tool that is analogous to the Makefile process and provides support for common operations such as invoking the Java compiler, the Java archiver tool used for creating libraries, the Javadoc tool used for creating API documentation from commented source code, as well as file copying, and directory creation/deletion operations.

The following steps are followed to create a new application specific portal called **ChemPortal** using the GPDK:

- Compile core GPDK service beans and portal objects.
When GPDK is compiled for the first time using the ANT build script e.g. `build all`, all of the necessary classes are compiled and placed into the GPDK Java Archive Repository (JAR) file.
- Invoke the ANT build tool to create a new portal with the desired name e.g. `build new ChemPortal`

By specifying a new project target, e.g. `ChemPortal`, when invoking the GPDK build script, a new portal sub project, `ChemPortal`, is created containing the pre-processed template portal pages. The template portal pages are composed of functional Java source code, template action page classes and view pages in JSP in addition to Javadoc generated API documentation for the generated source code. File pre-processing capabilities that allow the GPDK to generate a complete template portal project is made possible by the substitution capabilities provided by the ANT build tool. Ant pre-processes the template file names as well as the contents of the file to provide project specific Java code. For example, the template GPDK Servlet (GS) is pre-processed to create a `ChemPortalServlet`. The Portal Engine (PE) template is pre-processed to create a `ChemPortal` object that uses the core GPDK Portal object for initialization and shutdown routines. Similarly, the GPDK User Profile object is sub-classed to create a `ChemPortal User Profile` object that can be edited to contain data pertinent to the users of the `ChemPortal`.

The specified project namespace also provides a packaging namespace for portal source code. Core GPDK objects and beans are packaged under the `org.gpdk` namespace as per the Sun source code packaging convention and the pre-processed template project code is packaged under a new project namespace e.g. `org.chemportal`. This logical separation of core GPDK service beans and project specific classes allows for the creation and deployment of multiple portal projects to the same application server that can make use of the standard set of GPDK service beans.

- Modify and create new Action Pages (AP) and JSP view pages for project specific portal operations

As new portal operations are needed, `ChemPortal` developers would create a new Action Page object that implements the Page interface. The Page interface defines only one method, `execute`, that must be filled in by the developer. A correspond view page written as a JSP would need to be developed to return a display of the operation results to the client. As an example, imagine a new capability that allows `ChemPortal` users to run a GAMESS [10] simulation with a particular set of parameters. (GAMESS is a scientific



code used for performing ab-initio chemistry simulations). The ChemPortal developer would create a `RunGAMESSPage` class that would use the GSDK Job Submission beans to initiate a GAMESS simulation with the appropriate run parameters. Furthermore, the GSDK file transfer beans could be used to move any generated output files to a storage server for permanent archiving. Finally, the developer would create a `GAMESSOutput` JSP that would be used to provide information on a completed run to the end user.

- Deploy new portal project to application server
Simply invoking the ANT build script e.g. `build.pl all` will package all of the web application files, HTML, java classes and JSP in the form of a web application repository (WAR) file and deployed to the Tomcat application server, where it can be accessed by a web browser.

The steps presented above allow an application developer to create a fully functional portal complete with template code in a matter of minutes. The build environment provides easy to use compilation and deployment capabilities, eliminating the need for additional Makefiles or project specific build scripts. In fact, most of the work done in developing a new portal is in the action pages and JSPs that are project specific.

6.1. Application Portals Based on GSDK

The Grid Portal Development Kit has proven successful in the creation of application specific portals under development by other research groups. The following list briefly describes various ongoing portal projects that have been developed using the Grid Portal Development Kit framework: The NASA Launchpad [3] user portal seeks to provide web based access to users of the NASA Information Power Grid (IPG) [31]. Launchpad is based entirely on GSDK and takes advantage of the GSDK service beans to allow IPG users access to high-performance computer resources and IPG Grid Information Services. The NASA Nebula portal provides a web-based interface to the Nebula simulation code. Scientists use the Nebula code to study the atomic composition of interstellar clouds. The Nebula portal allows researchers to create an appropriate list of input parameters to the Nebula simulation and submit their job to NASA IPG resources. The Astrophysics Simulation Collaboratory Portal (ASC) [38] is an application portal developed to provide astrophysics researchers web based access to the Cactus computational toolkit. Cactus is a framework for solving various wave equations with an emphasis on astrophysical simulations of colliding neutron stars and black holes. The ASC portal allows users to remotely checkout and compile Cactus with specialized options as well as submit various Cactus simulations to a distributed set of HPC resources. The NCSA portal development effort is focused on the development of several application portals for computational chemistry and particle physics simulations. GridGaussian, a computational chemistry portal allows users to provide input files to the popular Gaussian chemistry package and easily submit simulations via the web. Similarly, the MIMD Lattice Computation (MILC) portal is aimed at providing particle physicists access to a popular community code used to understand elementary particle interactions. Both portals rely on GSDK for job submission and file staging capabilities.



7. Related Work

While the GPDK demo portal is very similar to many other portal projects, the ability to develop new portals using the GPDK framework offers a new capability to the scientific community and is one of the first of its kind. As a Software Development Kit (SDK), the GPDK is most similar to the GridPort [37] Toolkit developed by the San Diego Supercomputer Center (SDSC) to facilitate the development of application specific portals. The GridPort toolkit is implemented in Perl and makes use of the existing HotPage [36] technology for providing access to Grid services. GridPort supports many of the same Grid services as GPDK including the Myproxy service, GRAM, and GIS.

Application portal development using GridPort is enabled in two ways: The first approach requires that Globus software tools be installed because the GridPort scripts wrap the C Globus command line tools in the form of Perl CGI scripts. The second method of developing a portal using GridPort does not require Globus, but relies on the CGI scripts that have been configured to use a primary GridPort portal as a proxy for access to GridPort Services such as authentication, job submission, file transfer, etc. In the second approach, an application scientist can quickly deploy a web server configured with a set of GridPort CGI scripts to perform very generic portal operations. However, the ease of deployment comes at a cost of portal customizability. Because the HTTP response from the proxy GridPort server contains the HTML to be displayed, the graphical interface displayed to portal users is the same as the base GridPort server. In addition, both logic and design are encapsulated in the GridPort Perl CGI scripts making code maintainability and customization more difficult over the approach taken by GPDK which separates the graphical interface provided by template JSP pages and the logic contained within the GPDK beans.

8. Conclusion and Future Work

The architecture of the GPDK is highly modular and provides easy re-use of Grid components and a useful framework for the development and deployment of application specific portals. Based on the Model View Controller design paradigm, the GPDK has proven to be an extensible and flexible software package in use by several other portal building efforts. The GPDK makes use of commodity technologies including the open-source servlet container Tomcat and the Apache web server. The Grid Portal Development Kit makes use of the Java Commodity Grid (CoG) toolkit for its pure Java implementation of client side Globus Grid services as well as other widely available, commodity Java libraries. Future work on the Grid Portal Development Kit involves development in the following areas:

- Improve and enhance the functionality of GPDK service beans to support emerging Grid services and develop new portal capabilities. The GPDK information service beans will be enhanced to support the improved and secure MDS-2 infrastructure as it becomes widely deployed. Other future developments include the integration of a secure database to store user profiles rather than maintaining them on the portal. Additional application information may also be stored in the database or in the information servers, for access



by the GPKD information service beans. Task composition from a portal has become increasingly important as portal operations become more complex. Portal users would like to specify Grid or portal operations as tasks and be able to combine tasks together to create a work flow system for an entire calculation involving staging data, running a simulation and migrating output data to a storage system. Emerging Grid monitoring and event technologies [22] will provide useful information on network and computer utilization and performance allowing portal users and scheduling software alike to make better informed resource scheduling decisions. As these technologies become widely deployed, components for accessing these services will be developed in the GPKD.

- Investigate and prototype new distributed computing technologies including web services. Web services have become increasingly important in the enterprise community and many new standards and implementations are emerging from Sun, Microsoft and IBM. The Web Services Definition Language (WSDL) [25] permits services to be defined by a standard interface and registered and discovered using the Universal Description, Discovery and Integration (UDDI) [26] specification. The Simple Object Access Protocol (SOAP) [27] provides a standard for communicating structured information using XML. GPKD client web services beans will be developed that are capable of exchanging SOAP messages with Grid enabled web services as they become widespread.

ACKNOWLEDGEMENTS

We are grateful to many colleagues for discussions on portal development and working with early incarnations of GPKD to improve on its usefulness and robustness. In particular we wish to thank Jarek Gawor, Gregor Laszewski, Nell Rehn, George Myers, Mark Wallace, Farah Hasnat, Yinsyi Hung, John Lehman, Jeffrey Becker, Michael Russell, Shawn Hampton, Scott Koranda and John White.

REFERENCES

1. Netscape Directory and LDAP Developer Central.
<http://developer.netscape.com/tech/directory/index.html> [22 Nov 2001].
2. The Jakarta ANT Project
<http://jakarta.apache.org/ant/index.html> [22 Nov 2001].
3. The NASA Launchpad User Portal
<http://www.ipg.nasa.gov/launchpad> [22 Nov 2001].
4. The Astrophysics Simulation Collaboratory
<http://www.ascportal.org> [22 Nov 2001].
5. Java Servlet 2.3 and Java Server Pages 1.2 Specifications
<http://java.sun.com/products/servlets> [22 November 2001].
6. Apache Webserver Project
<http://www.apache.org> [22 November 2001].
7. ModSSL
<http://www.modssl.org> [22 November 2001].
8. Jakarta Apache Project
<http://jakarta.apache.org> [22 November 2001].
9. Tomcat open-source servlet container
<http://jakarta.apache.org/tomcat> [22 November 2001].



10. The General Atomic and Molecular Electronic Structure System
<http://www.msg.ameslab.gov/GAMESS/GAMESS.html> [22 November 2001].
11. IAIK Security Libraries
<http://jcewww.iaik.at/> [22 November 2001].
12. Sun Java Secure Sockets Extension
<http://java.sun.com/products/jsse> [22 November 2001].
13. Allcock W, Bester J, Bresnahan J, Chervenak A, Liming L, Tuecke S. GridFTP: Protocol Extensions to FTP for the Grid. *Grid Forum Working Draft* March 2001; <http://www.gridforum.org>
14. Allen G, Russell M, Novotny J, Daues G, Shalf J. The Astrophysics Simulation Collaboratory: A Science Portal Enabling Community Software Development. *Proc. 10th IEEE Symp. On High Performance Distributed Computing* 2001;
15. Czajkowski K, Fitzgerald S, Foster I, Kesselman C. Grid Information Services for Distributed Resource Sharing. *Proc. 10th IEEE Symp. On High Performance Distributed Computing* 2001;
16. Czajkowski K, Foster I, Karonis N, Kesselman C, Martin S, Smith W, Tuecke S. A Resource Management Architecture for Metacomputing Systems. *Proc. IPPS/SPDP '98 Workshop on Job Scheduling Strategies for Parallel Processing* 1998;
17. Foster I, Kesselman C (editors) *The Grid: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann, 1998.
18. Foster I, Kesselman C. Globus: A Metacomputing Infrastructure Toolkit. *International Journal of Supercomputing Applications* 1997;
19. Grid Forum Web Site
<http://www.gridforum.org> [22 November 2001].
20. Globus Web Site
<http://www.globus.org> [22 November 2001].
21. GSI Software Information
<http://www.globus.org/security> [22 November 2001].
22. Grid Monitoring Architecture
<http://www.didc.lbl.gov/GGF-PERF/GMA-WG/> [22 November 2001].
23. The GPDK Project Page
<http://www.itg-lbl.gov/Grid/projects/GPDK/index.html> [22 November 2001].
24. The JMyproxy client
<ftp://ftp.george.lbl.gov/pub/globus/jmyproxy.tar.gz> [22 Nov 2001].
25. Web Services Definition Language
<http://www.w3.org/TR/wsd/> [22 Nov 2001].
26. Universal Description, Discovery and Integration
<http://www.uddi.org> [22 Nov 2001].
27. Simple Object Access Protocol
<http://www.w3.org/TR/SOAP/> [22 Nov 2001].
28. Foster I, Karonis N, Kesselman C, Koenig G and Tuecke, S. A Secure Communications Infrastructure for High-Performance Distributed Computing. *Proc. 6th IEEE Symp. on High Performance Distributed Computing* 1997;
29. Gamma E, Helm R, Johnson R, Vlissides J. *Design Patterns: Elements of Reusable Object Oriented Software*. Addison-Wesley, 1995.
30. Akarsu E, Fox G, Haupt T, Youn C. The Gateway System: Uniform Web Based Access to Remote Resources *IEEE Concurrency: Practice and Experience* July 2000;
31. Johnston W. E., Gannon D, and Nitzberg, B. Grids as Production Computing Environments: The Engineering Aspects of NASA's Information Power Grid. *Proc. 8th IEEE Symp. On High Performance Distributed Computing* 1999;
32. Laszewski G., Foster I, Gawor J. CoG Kits: A Bridge Between Commodity Distributed Computing and High-Performance Grids *Proceedings of the ACM Java Grande Conference* 2000;
33. Novotny J, Tuecke S, Welch V. An Online Credential Repository for the Grid: MyProxy. *Proc. 10th IEEE Symp. On High Performance Distributed Computing* 2001;
34. Pancerella C. M., Rahn L. A, and Yang C. L. The Diesel Combustion Collaboratory: Combustion Researchers Collaborating over the Internet. *Proc. of IEEE Conference on High Performance Computing and Networking* November 1999;
35. Romberg M. The UNICORE Architecture *Proc. of the 8th IEEE Intl. Symp. on High Perf. Dist. Comp* 1999;



-
36. Boisseau J, Mock S, Thomas M. Development of Web Toolkits for Computational Science Portals: The NPACI HotPage. *Proc. of the 9th IEEE Intl. Symp. on High Perf. Dist. Comp* 2000;
 37. Boisseau J, Dahan M, Mock S, Mueller K, Sutton D, Thomas M. The GridPort Toolkit: A System for Building Grid Portals. *Proc. of the 10th IEEE Intl. Symp. on High Perf. Dist. Comp* 2001;
 38. Allen G, Daus G, Foster I, Laszewski G, Novotny J, Russell M, Seidel E, Shalf J. The Astrophysics Simulation Collaboratory Portal: A Science Portal Enabling Community Software Development. *Proc. of the 10th IEEE Intl. Symp. on High Perf. Dist. Comp* 2001;