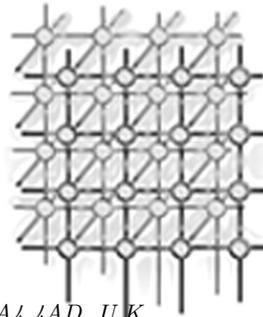


Developing Portals/Portlets using Enterprise JavaBeans for Grid Users



X. Yang^{*,†} and R.J. Allan

CCLRC e-Science Centre, CCLRC Daresbury Laboratory, Warrington, WA4 4AD, U.K.

SUMMARY

A web based portal acting as a presentation layer to the Grid allows us to provide a transparent means of access to Grid resources. Our previous experiences in developing Grid portals have shown that there is a need to draw a clear high-level picture of the architecture. This could be defined using the J2EE component-oriented architecture, which enables a clear separation between the presentation, business logic and data layers. Additional benefits come with J2EE 1.4 making it possible to build up advanced service-based Grid portals by exposing stateless session beans as web services. In this paper, lessons learnt from our prototype portal/portlet development using Enterprise JavaBeans (EJB) are reported.

KEY WORDS: portal; portlet; enterprise javabeans; web services; Grid

1. INTRODUCTION

Web based Grid portals such as the Telescience Portal [1], the HotPage Grid Computing Portal [2] and the BioSimGrid Web Portal [3] are prevailing in Grid communities because of their support for transparent access to resources. Portals are essentially web pages that provide a centralised access point and tools which can vary from a simple application such as an e-mail client or a web search engine to a complex Grid application. Today, portlets are widely adopted in constructing portals. These portlets are built with extensive usage of technologies including Java Servlet and JavaServer Pages (JSP) providing the ability to create dynamic web contents. The MVC (Model-View-Controller) programming pattern is widely utilised in portlet development. Both Struts and JavaServer Faces (JSF) are now supported by different portal frameworks allowing more sophisticated design patterns. For example, GridSphere [4] supports

*Correspondence to: CCLRC e-Science Centre, CCLRC Daresbury Laboratory, Warrington, WA4 4AD, U.K.

†E-mail: x.yang@dl.ac.uk



portlet development using JSF and “portlet-ising” Struts applications using the Apache Portals Bridges [5].

To solve the interoperability issues in portal/portlet development, two standards – WSRP (Web Services for Remote Portlets) 1.0 [6] and JSR 168 (Java Portlet Specification) [7] were proposed by OASIS and JCP respectively in 2003. As mentioned in [8], whilst WSRP is a universal protocol for communication between different portals, JSR 168 is a set of Java APIs that standardise communications between Java portlets and a portlet container. With these two specifications, portlet developers can concentrate on developing standard JSR 168 portlets, while portal frameworks would provide the WSRP support for publication of portlets (as producers) and consumption of remote portlets (as consumers), i.e., portlets inside remote portlet containers. The JSR 168 specification is adopted by most of the portal frameworks today but we have found the WSRP support to be still immature [9]. While some of the portal frameworks provide WSRP consumers to consume remote portlets, almost none of their WSRP producers can be consumed by another portal framework without issue. In addition to being published as remote portlets within a normal portal framework which has WSRP producer support, portlets can be published through 3rd-party WSRP producers like WSRP4J [10], which in turn makes use of Pluto [11] as its portlet container. As reported in [9], neither producers nor consumers are however fully functional. In this paper, therefore, only the use of JSR 168 is considered.

No matter how complex portals become, they are normally treated as a presentation layer to business logic which manages special features like customisation and single sign-on (SSO). With development experience in several portal projects such as the UK NGS (National Grid Services) Portal [12] and the Cambridge CFD Grid Portal [13], we believe that a clear separation between presentation and business logic layers is best practice for portal/portlet development. Although simple portals can combine presentation and business logic for convenience, it is hard to re-use the code. Our experience shows that for more sophisticated portals such an approach is neither convenient during development nor later for maintenance. Business logic can be implemented either as some simple Java classes or complex J2EE components such as Enterprise JavaBeans (EJB). Theoretically, selection of technology is governed by the requirements and complexity of business logic. Today, Grid portals are becoming more and more complex in their functionalities and need persistence, security, transaction control, load balancing and much more. We therefore decided to investigate implementing business logic and data layers as EJB.

J2EE developers often consider the use of EJB to be a heavyweight technique owing to their complexity and inefficiency. As a widely adopted industrial standard however, J2EE/ EJB is quite mature and therefore we have explored how the J2EE architecture can benefit us in portal/portlet development. The web service support in the EJB 2.1 specification extends the flexibility of EJB which are now no longer limited to RMI-IIOP. The new EJB 3.0 specification (JSR 220) tries to solve the complexity of current EJB techniques by introducing the programming annotation facility of Java 5 and simplifying of entity persistence through the Java Persistence API, which brings ideas from leading persistence frameworks and APIs such as Hibernate [14] and Oracle TopLink [15]. These new features of EJB plus existing values of mature J2EE application servers such as scalability, transaction and security support make it a promising technology for development of service-based portals.



In this paper, we will first describe the J2EE based component-oriented architecture (COA) with focus on using Enterprise JavaBeans in our prototype portal/portlet development work. Then we describe how to convert a COA to a service-oriented architecture (SOA). This is followed by a description of our initial work on developing Grid portlets based on EJB. Some related work on advanced serviced-oriented Grid portals is also presented followed by conclusion remarks.

2. COMPONENT-ORIENTED ARCHITECTURE AND SERVICE-ORIENTED ARCHITECTURE

2.1. Component-Oriented Architecture (COA)

The J2EE architecture is an effective component-oriented architecture. The EJB technology, one of the key technologies in J2EE, provides the server-side component architecture for J2EE. EJB model both business and data layers by defining session beans and entity beans. Without worrying about how to present information, session beans focus on how to manage data for users. This could involve complex algorithms, business procedures, etc. Entity beans manage objects in persistent storage and talk to them to either fetch data from or write data to relational databases. Especially with container-managed persistence, the heavy duty database operations are now the task of the application server rather than programmers. Besides session and entity beans, the use of a third-type of EJB, message-driven beans, offers a robust and scalable enterprise messaging solution to leverage existing enterprise applications.

An extra presentation layer should normally be added to COA to provide the user interface (UI), see Fig. 1. The nature of EJB makes it possible to have multiple UI clients such as a desktop program or a web application. Today, web interfaces are widely adopted for e-Business and e-Science. J2EE provides Java Servlet, JavaServer Pages and JavaServer Faces to generate dynamic web contents for rendering information in web browsers. As mentioned previously, by extending the servlet specification, web pages managed by portals can now be composed of a set of web components called portlets.

2.2. Service-Oriented Architecture (SOA)

Web services provide a useful way for building up a message-based service-oriented architecture. The core idea of SOA is loose coupling among interacting software agents. A service on the service provider side is usually responsible for executing a set of data-centric tasks, which are wrapped with only some simple APIs published. Because of this publishing mechanism a web service based SOA is suitable for constructing large-scale applications that integrate heterogeneous distributed services. Although SOA is commonly realised through web services, it can also be done using other techniques such as CORBA (Common Object Request Broker Architecture).

COA and SOA are in fact complementary architectures. A COA-based system can be converted to a SOA-based system in a straightforward way by exposing EJB as web services. The EJB 2.1 specification brings support for this through exposing stateless session beans as

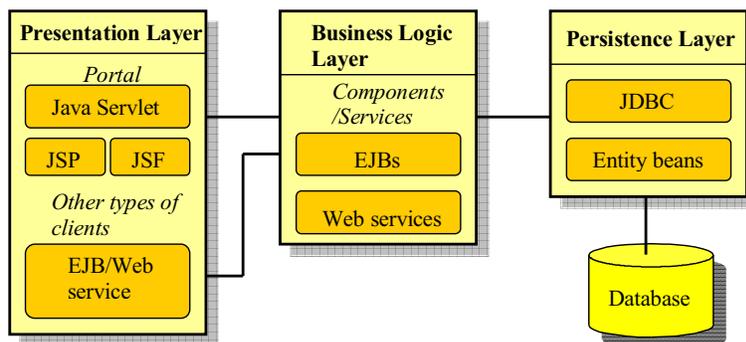


Figure 1. COA/SOA with portal acting as the presentation layer.

web services. Thus when we talk about COA in this paper, it can also imply SOA. By adopting the loosely coupled COA, components can be platform independent and can be easily shared, e.g. between the Java and Microsoft .NET programs. By re-using EJB or remote services, portal/portlet developers can focus on how to present data rather than writing and placing complex business logic within portlets. This greatly reduces the burden on portal developers in deploying large complex portal systems. The drawback here is that besides a servlet/portlet container, an extra EJB container is needed which makes the overall system more complex. EJB development and testing are also much more difficult compared with POJOs (Plain Old Java Objects). The benefit however is that once deployed the services can be re-utilised in many different portals to support a diverse range of projects and end-user communities.

3. PROTOTYPE PORTAL EXAMPLE

3.1. NGS Portal

The example we are going to discuss is based on our early work on portal development for the UK National Grid Service (NGS) [12]. NGS is the UK's core production computational and data Grid which runs Globus Toolkit 2 (GT2) and other services including SRB [16]. StringBeans [17], an open-source JSR 168 compliant portal framework, was adopted as the choice for the NGS Portal mainly because of its powerful tools for administrative tasks. It has been customised to accept authentication through the UK NGS MyProxy [18] server which is maintained separately from the portal server. As StringBeans supports the Java Authentication and Authorisation Service (JAAS) pattern, a MyProxy login module was written to realise our customised authentication through MyProxy. With this module, end-users can now use Grid tools after successfully logging into the portal, since their Grid proxy credentials are retrieved



automatically during the login procedure. Some JSR 168 portlets were developed to support the generic functionalities provided by GT2 with Grid proxy management, job submission, job monitoring, file transfer (using GridFTP [19]) and LDAP/MDS (Metacomputing Directory Service) [20] browsing. SRB (Storage Broker Broker) is also supported. Currently, data providers are included in the portlets with database access based on Hibernate [14].

3.2. Example portal based on EJB

In this section we will not cover all the portlets mentioned above but focus on MDS/ LDAP query, proxy manager and job submission. The basic idea is to keep our Grid portlets as simple as possible, providing easily adaptable presentation layer to the Grid applications. The business logic has been removed from the portlets and is implemented in session beans instead. In the demonstration portal application, users log in using their portal usernames and passwords, that is, the login procedure is managed by the portal. We did not apply the MyProxy login module in this test, so there is no modification of the portal framework itself. Authenticated users are given access to different portlets by the portal according to their roles as specified by the portal administrator.

A set of session beans and entity beans have been written for portlets to consume. These include two container-managed persistence (CMP) entity beans – each of them maps to a database table. The first table is designed to maintain user profiles in particular proxy credentials and the second one is responsible for recording metadata of jobs submitted by end-users. As we can see from Fig. 2, NgsUsers and NgsUserJobs are mapped to two tables, *ngs_users* and *ngs_user_jobs* in the left corner. These two entity beans are wrapped by two session beans for external access. An additional session bean, the LdapQuery bean, has been developed for retrieving information such as CPU type and memory size from the NGS MDS server. Similarly to portlets in the NGS Portal, these Enterprise JavaBeans are built up by using Java CoG 1.2 [21].

Based on the EJB mentioned above, we have developed three portlets for demonstration purposes. A LDAP Browser portlet for MDS query, a Proxy Manager portlet for retrieving/renewing Grid proxy credentials from the NGS MyProxy server and a Job Submission portlet for submitting jobs to a remote Globus GRAM Gatekeeper. These three portlets were developed from the portlets used in the NGS Portal with only minor modifications in replacing the built-in business logic and persistence parts with EJB client code. The EJB and portlets together with the StringBeans portlet container have all been successfully deployed in the JBoss Application Server [22]. Fig. 3 gives a screenshot of these three portlets running inside StringBeans.

The LdapQuery session bean has also been successfully exposed as a web service with the following steps applied:

- (i) define an endpoint interface with all operations to be published in the web service. The endpoint also needs to be defined in the *ejb-jar.xml* configuration file;
- (ii) create deployment descriptors including *webservices.xml* (which defines a set of web services that are to be deployed in a J2EE enabled container) and a JAX-RPC mapping

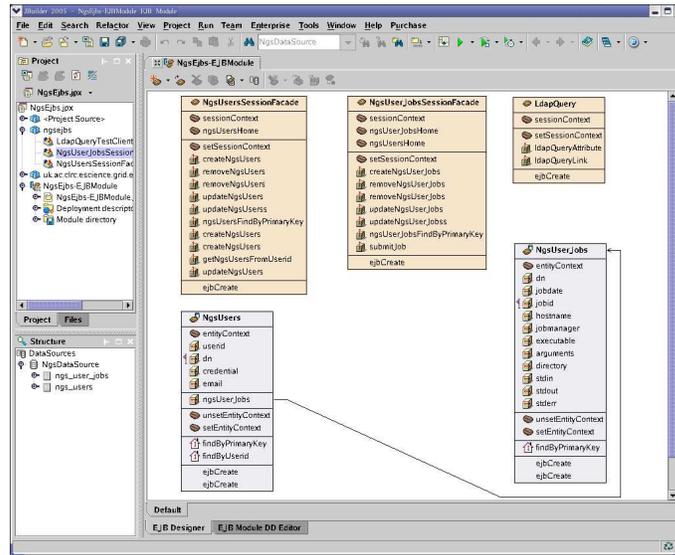


Figure 2. EJB developed for our example portal.

Proxy Manager Portlet

Your available GSI proxy credential is shown below.

Subject: C=UK,O=eScience,OU=CLRC,L=DL,CN=xaobao yang,CN=proxy,CN=proxy,CN=proxy

Issuer: C=UK,O=eScience,OU=CLRC,L=DL,CN=xaobao yang,CN=proxy,CN=proxy

Bits: 512

Time: 1 hour(s) 59 minute(s) 59 second(s)

Left: 1 hour(s) 59 minute(s) 59 second(s)

Refresh **Renew**

Job Submission Portlet

Host Name: jsh@mpuhome.ac.uk

Job Manager: jsh

Executable: jsh

Arguments: j

Working Directory:

Standard Input File:

Standard Output File: js.out

Standard Error File:

Job Type: Batch Interactive

Launch **Default**

GRAM Job Result: Job submitted successfully

LDAP Browser Portlet - Browse

LDAP Server URL: ngsinfo.grid-support.lac.uk:2135

Current Node: Mds-Host-Im-grid-data.rl.ac.uk.Mds-Vo-name=ngsinfo, o=grid

Attribute	Value
Mds-Cpu-speedMhz	3066
Mds-Memory-Ram-Total-freeMB	2158
Mds-Fs-freeMB	11504
Mds-Fs-freeMB	1293377
Mds-Fs-freeMB	1596
Mds-Fs-freeMB	1718
Mds-Fs-freeMB	1992
Mds-Fs-freeMB	2002
Mds-Fs-freeMB	453247
Mds-Fs-freeMB	5299
Mds-Fs-freeMB	60726
Mds-Fs-freeMB	74014
Mds-Fs-freeMB	856
Mds-Fs-freeMB	920
Mds-Cpu-Free-SpinX100	378
Mds-Net-Total-count	3
Mds-validfrom	20050728150350Z
Mds-Cpu-Total-count	4
Mds-Memory-Vm-sizeMB	3843
Mds-Cpu-vendor	GenuineIntel
Mds-Net-name	eth0
Mds-Net-name	eth1

Figure 3. Screenshot of portlets running in StringBeans.



- file (which defines classes used for serialisation/de-serialisation of custom data types) with vendor-dependent configuration files;
- (iii) re-compile and package everything including EJB and deploy them.

The core LDAP query functions are now available for non-HTTP clients to access which makes the whole system much more flexible to suite different scenarios. This demonstrates the flexibility of the J2EE architecture since from version 1.4 it begins to support web service and the COA can simply become a SOA.

3.3. Hibernate vs. entity bean

From the programmers' point of view, it is always easier to understand and more efficient to handle objects rather than direct relational database operations. As mentioned in Section 3.1, Hibernate [14] is utilised in the current release of the NGS Portal as the data persistence service. Hibernate, one of the O/R mapping technologies, is widely adopted for developing persistent classes in Java. It can either store Java objects in a relational database or provide an object-oriented view of existing relational data. Similar to EJB which has EJB Query Language (EJB-QL), Hibernate has its own portable SQL extension (HQL). Both techniques are attempts to simplify the data layer management for software developers. Unlike entity beans, Hibernate does not require a container but only some libraries to be included in an application. It is often argued in the Java community that CMP entity beans have disadvantages compared to O/R mapping technologies such as Hibernate:

- (i) persistence technology within a container normally involves other container-based services which makes it complex, for instance, difficult for testing;
- (ii) entity beans are not portable between application servers;
- (iii) without an IDE (Integrated Development Environment), CMP entity bean development can involve a lot of work.

Since the proposed EJB 3.0 specification (JSR 220) could make it a lot easier to developing EJB, Hibernate is beginning to support the EJB 3.0 persistence standardisation effort. In JBoss, it is now used as a container-managed persistence service.

For the persistence service, there is no longer much difference between entity beans and Hibernate, especially in EJB 3.0. Persistence is also only part of the EJB specification. The benefit of adopting EJB lies mainly in other types of EJB (session bean and message-driven bean) and additional functions provided by the application server. In fact, it is possible to develop J2EE applications without even using EJB by adopting frameworks like Spring [23], but this is out of the scope of this paper.

4. RELATED WORK

Many Grid portals are nowadays *service* based. Besides providing traditional functions, advanced service based Grid portals normally provide part or even all their functionalities as web services which can be consumed outside of the traditional portals/portlets in a SOA.



GridPort 3 [24] adopts a clear 3-tier architecture which includes a J2EE web server tier, a J2EE application server tier and a database and service tier. This is very similar to the J2EE 3-layer architecture illustrated in Fig. 1. In GridPort 4 the architecture has been re-designed with three layers called portal layer, service layer and resource layer. Compared to the COA-based one, this places emphasis on the importance of *service* and *resource*. For example, Grid related functions such as GridFTP and GRAM can be treated as generic services while database and HPC are included in the resource layer. Based on SOA, GridPort 4 thus provides users with a simple, easy-to-install software package for creating portals to access Grid infrastructure via a web interface.

NEESit [25] links earthquake researchers together across the USA with computing resources and engineering facilities. It provides a virtual research laboratory for data management, collaboration, telepresence and simulation. In order to link researchers effectively, web-based portals such as NEEScentral and NEESport [26] have been developed. By aggregating and coordinating lower-level services provided by the Globus Toolkit, NEESport provides high-level middle-tier services for seamless access to remote resources. In NEESport again the concepts *service* and *resource* are explicitly emphasised.

Another U.S. based example is the GEON (Geoscience Network) cyber-infrastructure project [27] which aims at creating an infrastructure to enable virtual laboratory for geoscience researchers. A service-oriented architecture is also adopted with a portal developed for end users to access a set of portlets.

5. CONCLUSIONS

Current portals developed for Grid communities normally bind presentation, business logic and data sources within portlets. In this paper, we have discussed applying the J2EE pattern-based techniques especially the Enterprise JavaBeans to model business logic and persistence layers. EJB can act as information/data sources for portals/portlets. The benefit of adopting mature J2EE techniques in portlet development brings a lot of added value such as reusability, scalability, security and transaction support which need further investigation. With data providers extracted from portlets as EJB, the structure of our portal system becomes very clear with a focus on presentation only. By using container-managed persistent entity beans, the J2EE application server handles data with automatic transaction support. Components written as stateless session beans can be easily converted to web services enabling the building up of a suite of services for re-use in advanced service-oriented Grid portals serving divers user communities. Alternative (non-web) interfaces can also be provided.

Whilst a basic portal based on EJB has been demonstrated successfully, this is only the starting point of adopting J2EE technology. As there have been arguments against using J2EE and EJB on grounds of complexity, this study has enabled us to identify the benefits and difficulties in the context of our own work. We are now beginning to plan how to use J2EE technologies to build up much more complex systems than the prototype portal example. The final system envisaged will permit ongoing development and maintenance of business logic as a single set of code which can be used in several portal interfaces and also exposed as web services to be consumed by desktop applications requiring access to Grid resources such as the



GROWL Virtual Research Environment (VRE) [28] project. This may involve investigation in Java EE 5 and the Spring framework.

On the other hand, portals can be built up on top of existing legacy systems. If such a system happens to be constructed using EJB, as shown in this paper, these systems can then be easily integrated in web portals.

ACKNOWLEDGEMENTS

We thank Asif Akram for his help in EJB development and anonymous reviewers for their insightful comments helped to improve this paper. This work has been undertaken at the CCLRC e-Science Centre supported by UK JISC (The Joint Information Systems Committee).

REFERENCES

1. Peltier S.T., Lin A.W., Lee D., Mock S., Lamont S., Ellisman M.H. The Telescience Portal for Advanced Tomography Applications. In *Journal of Parallel and Distributed Computing*, 2003; **63**(5):539-550.
2. HotPage Grid Computing Portal. <https://hotpage.npaci.edu/> [1 December 2005].
3. Wu B., Dovey M., Ng M.H., Tai K., Murdock S., Fangohr H., Hohnston S., Jeffreys P., Cox S., Essex J.W., Sansom S.P. A Web/Grid Portal Implementation of BioSimGrid: a Biomolecular Simulation Database. In *Journal of Digital Information Management*, 2004; **2**(2):74-78.
4. GridSphere Portal Framework. <http://www.Gridsphere.org/> [14 July 2006].
5. Portals Bridges. <http://portals.apache.org/bridges/> [14 July 2006].
6. WSRP Specification 1.0 by OASIS. <http://www.oasis-open.org/committees/download.php/3343/oasis-200304-wsrp-specification-1.0.pdf> [1 December 2005].
7. JSR 168: Portal specification. <http://www.jcp.org/en/jsr/detail?id=168> [1 December 2005].
8. Vickers S. Portal Standards the Answer to Portal Interoperability? In *Java Developer's Journal*, 2005; **10**(1):18-20.
9. Yang X., Wang X.D., Allan R.J. WSRP Support Investigation of Selected Open-Source Portal Frameworks. *Concurrency Computat.: Pract. and Exper.*, accepted.
10. WSRP4J. <http://ws.apache.org/wsrp4j/> [1 December 2005].
11. Pluto. <http://portals.apache.org/pluto/> [1 December 2005].
12. Yang X., Chohan D., Wang X.D., Allan R.J. A Web Portal for the National Grid Service. In *UK e-Science AHM 2005*, Nottingham, UK, 2005; available on CDROM.
13. Yang X., Hayes M., Jenkins K., Cant S. The Cambridge CFD Grid Portal for Large-Scale Distributed CFD Applications. In *Proc. ICSS2004*, eds. M. Buak etc al., Springer-Verlag, LNCS 3036, pp. 478-481, 2004.
14. Hibernate. <http://www.hibernate.org/> [1 December 2005].
15. Oracle TopLink. <http://www.oracle.com/technology/products/ias/toplink/index.html> [20 July 2006].
16. Baru C., Moore R., Rajasekar A., Wan M. The SDSC Storage Resource Broker. In *Proc. CASCON'98*, Toronto, Canada, 1998.
17. StringBeans. <http://www.nabh.com/projects/sbportal/> [1 December 2005].
18. Novotny J., Tuecke S., Welch V. An Online Credential Repository for the Grid: MyProxy. In *Proc. HPDC-10*, IEEE Press, 2001, pp. 104-111.
19. Allcock W., Bresnahan J., Kettimuthu R., Link M., Dumitrescu C., Raicu I., Foster I. The Globus striped GridFTP Framework and Server. In *Proc. SC05*, to be published.
20. Fitzgerald S., Foster I., Kesselman C., von Laszewski G., Smith W., Tuecke S. A Directory Service for Configuring High-Performance Distributed Computations. In *Proc. HPDC-6*, 1997; pp. 365-375.
21. von Laszewski G., Foster I., Gawor J., Lane P. A Java Commodity Grid Kit, In *Concurrency Computat.: Pract. and Exper.*, 2001; **13**(8-9):643-662.
22. JBoss Application Server. <http://www.jboss.com/products/jbossas> [1 December 2005].
23. Spring Framework. <http://www.springframework.org/> [1 December 2005].
24. The GridPort Toolkit. <http://Gridport.net/> [1 December 2005].



-
25. NEESit. <http://it.nees.org/> [1 December 2005].
 26. NEESport: Simulation Portal for NEESGrid. http://www.erc.msstate.edu/haupt/NEES_WEB/NEESport.html [1 December 2005].
 27. Cyberinfrastructure for the Geosciences. <http://www.geonGrid.org/> [1 December 2005].
 28. GROWL: Grid Resources on Workstation Library. <http://www.growl.org.uk/> [1 December 2005].