

Service Oriented grid Architecture for Geosciences Community

R.Fraser¹, T.Rankine² and R.Woodcock¹

(1) CSIRO Exploration and Mining and pmd*CRC, 26 Dick Perry Ave, Kensington, Western Australia 6151, Australia

(2) IVEC, 26 Dick Perry Ave, Kensington, Western Australia 6151, Australia

Ryan.Fraser@csiro.au

Abstract

Grid computing has many attributes that provide benefits to research communities. These include the ability to do existing research faster; enable collaborative environments; reduce costs and provide a pool of resources that creates an ability to conduct greater research. These have been identified by the Geosciences community as of significant interest and a plausible solution to some of the community's current linkage and development problems.

The Geosciences community for sometime now have often been restricted to "small" science, easy access to greater computational and data resources that would allow larger experiments to be conducted. The Grid poses a solution to this problem by providing researchers simple uniform access to a myriad of resources via a single-sign-on policy.

The Australian Partnership for Advanced Computing (APAC) Geosciences Grid project together with Solid Earth & Environment Grid (SEE Grid), aims at providing Computational Geoscientists with uniform access to some of the more common Geosciences-based computational codes on large computational and data resources available via the grid. To achieve this goal, an interoperable approach has been taken and web/grid services where used to establish the links to such resources.

The Grid's service oriented architecture enables the programming of specific self-contained services that simplified the access to computational and data resources for the user. Through the use of the Globus Toolkit 4, we have developed grid services that assist the common Geoscientist user to collaborate and conduct large-scale experiments including parameter sweeps.

Copyright (c)2007, Australian Computer Society, Inc. This paper appeared at the Australasian Symposium on Grid Computing and Research (AusGrid), Ballarat, Australia. Conferences in Research and Practice in Information Technology, Vol. 68. Editors, Ljiljana Brankovic, University of Newcastle, Paul Coddington, University of Adelaide, John F. Roddick, Flinders University, Chris Stekete, University of South Australia, Jim Warren, the University of Auckland, and Andrew Wendelborn, University of Adelaide. Reproduction for academic, not-for profit purposes permitted provided this text is included.

The deployment of these interoperable grid services have allowed us to develop multiple client programs in varying languages in an effort to give researchers the choice of tools to use to interact and conduct their research. We have also used existing client tools to connect to the. Such interfaces have been deployed to several Geoscientists and into the pmd*CRC.

Keywords: Service oriented Architecture, ws-rf, grid

1 Introduction

The grid provides a standardised method to access domain-independent computational and data services such as job management, data storage, monitoring and security. These attributes provide great benefits to research communities including the ability to do existing research faster, enable collaborative environments, reduce costs and the pooling of resources, allows for the ability to conduct greater research.

Computational Geosciences is concerned with the development of computational technologies for the geosciences community. The community looks at the construction of computational models of geology at all scales. Computer simulations are used to model the geological processes which contribute to the formation of different types of deposits. Visualisation of the results is also a core component of their workflow.

The main problems that the Geosciences Community experience are "linkage"; the interactions in the community with other persons, organisations and resource (computational and service), multiple inefficiencies exist here.

The Geosciences support project in the APAC grid programs along with SEE Grid information standards, has the goal of developing service interfaces to the communities more commonly used computational codes and algorithms (Woodcock & Fraser, 2006b) to aid in improving the "linkage" problem being experienced in the community. These service interfaces are built on top of domain-independent grid technologies and are made available at multiple sites on the APAC grid using a standardised interface.

This paper will discuss the motivation of the project and justify its intent. We will describe the architecture used to provide better linkage for Computational Geoscientists to computational and data resources. Finally we will describe the pmd*²CRC Desktop Modelling Toolkit (DMT) that makes use of the architecture.

2 Motivation behind the project

The motivation for this project is to improve the links within the Geosciences community to computational resources and services. Currently there are several inefficient processes within the typically workflow and typically occur between people, organisations and resources. Some examples include:

- there are multiple HPC resources available to users in Australia, but these are often extremely difficult to use due to differing site policies, queuing techniques and data requirements at each site.
- Data and Information relevant to workflows is scattered across multiple geological surveys and the mining companies. This hampers the gathering of input data. As a consequence investigators in the Geosciences community often use the “average” properties of observations, ignoring real world observational data due to the cost of data integration. The problem with this is that “average” properties can range by range by several orders of magnitude between geographic locations.

Furthermore, in the Geosciences field, like many other sciences, investigators and researchers often need high performance computing to explore greater depths in their science.

Why do the Geosciences community require interoperability? When trusted and sharing relationships are required between multiple arbitrary parties, across different domains and platforms, languages and environments- interoperability is a concern. If mechanisms for communication between boundaries are not defined and implemented so that they can be communicated with by anyone/thing, we will not achieve our goal of being able to collaborate effectively. Interoperability and the grid will enable the geosciences community to have access to a greater amount of resources to improve their science and collaboration.

Why are services and the grid the solution? To solve the problem of making the Geosciences community interoperable and being able to participate in “big” science, they need access to computing services that can be connected together as small highly cohesive components to achieve their desired science outcomes. Each service which is defined by the protocol(s) it uses and the behaviours it implements, completes a well-defined function in a workflow of achieving the end goal.

Grid technologies provide support for the sharing and coordinating of distributed resources which are managed by dynamic Virtual Organisations (VOs) (Burbeck, 2000). The grid provides the “best fit” solution for enabling the geosciences community access to HPC resources by simplifying the process to gain access and use the resources available via the grid. It also provides significant opportunities to improve linkage and chain services together to construct more and more complex workflows to meet demands of today’s researchers.

3 System Architecture

3.1 Overview of the architecture

The grid is an extensible, distributable, scalable service oriented architecture (SOA) system (Lin, 2005). SOA architecture for the grid provides methods for exposing services and allowing computers to talk to each other in a highly heterogeneous environment. Services on the grid may be individually purposeful (fully self-contained) or act as a component may contribute a useful workflow. Services communicate with clients by exchanging the messages that define them and also publish their capabilities

The benefits of a SOA architecture on the grid is that it provides loose coupling which results in the system being flexible, scalable, services are replaceable and adds fault tolerance (Srinivasan and Treadwell, 2005). For these reasons, a service oriented architecture for the Geosciences grid has been used and implemented in our system.

The following sections will describe the typical services that participate in the Geosciences grid system which enable users to submit the same job to the grid and that job could run at multiple sites.

3.2 Architecture Components

The choice of architecture for the geosciences grid project was based on the intended usage of the Globus Toolkit 4 (GT4 - <http://www.globus.org/toolkit/>), programming ws-rf (<http://www.globus.org/wsrf/>) based services. This architecture is not Geosciences specific – it may be applied as a solution for other communities; it is only presented here in the Geosciences context.

Each computational code that was deployed on a cluster on the APAC grid, a ws-rf service was implemented for it and deployed to the GT4 web container at the site where the code is installed. This service we title “ws-*”, where the “*” is replaced by the code name (for example – ws-FastFlow for the Fast Flow computational code). Figure 1 depicts the architecture implemented to solve the problem.

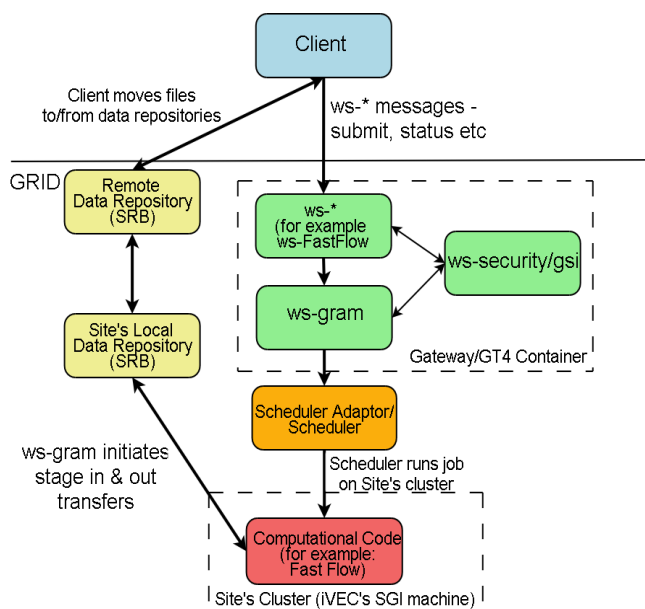


Figure 1. Geosciences Grid SOA architecture

This architecture was chosen because of its scalability and ease to replicate across the APAC grid. Further decisive factor was the ease of use this architecture bestowed to the client, a job could be submitted to any site on the grid and only the URL for the site would need changing and inputs and outputs could be staged from the single federated data store.

3.2.1 The Client Application

Client applications in this architecture can be implemented in any language that can create and ingest web service messages (for example – SOAP messaging). In our workflows, the clients typically act as either a user interface or alternatively – another service. Client applications, that interact with the "ws-*" have been implemented in python, c and java. A client program marshals data to (job input data) and from (results) from the Data Repository. It sends messages to the ws-* service such to submit jobs, terminate the job, get job's status, the standard output/error and retrieve the intermittent result sets.

3.2.2 The Service – "ws-*

The ws-* service is implemented as a GT4 service and deployed into the GT4 web container. The Service was coded using Java and the Globus java ws-core api [http://www.globus.org/toolkit/docs/4.0/common/javaws-core]. Essentially this Service is a leverage point for implementing intelligence for job submission and management. The Service is a place where all the Code-specific details can be sorted out on the users' behalf, for example:

- Identifying where to submit jobs
- Finding sites that support the requirements of a job

- Setting up paths and fleshing out the full job description document for submission to the Job Submission service (ws-gram)

The Service monitors jobs, acts upon messages and notifications from the client (submit job, terminate job, get job status etc) and other services (resource destruction, job status etc.) The Service is an integral component in developing service chained workflows and "smart" services.

3.2.3 GT4 and the Job Submission Service (ws-gram)

GT4 is a toolkit containing many services that have web service front-ends, including services for job submission (ws-gram), file transfer (Reliable File Transfer - rft and gridftp), security (ws-security & gsi) and delegation among many others. User Services are coded and deployed to the GT4 container, where each can interoperate in a secure environment.

The ws-gram service in the GT4 container is of particular interest. The ws-gram service is the job submission service in the GT4 toolkit and is used as a standardised job submission interface to clusters/sites on the grid. ws-* constructs a job submission file that is submitted to ws-*, ws-gram returns an endpoint reference to the job.

3.2.4 GT4 scheduler adaptor and the Local site scheduler

The GT4 scheduler adaptor converts the job script submitted by ws-gram into a script that can be submitted to the local site's scheduler. The adaptor then submits the converted script to the local scheduler which returns a reference to the job submitted.

3.2.5 Computational Code – Fast Flow (pmd*RT), eScript, CitcomS, Abaqus/Calculix, Underworld/Snark

Computational codes are installed on a site's cluster and are accessible on the cluster via modules which setup the user environment. The local scheduler manages and monitors the job running on the cluster and notifies the ws-gram with status updates, which then relays the notification to ws-* which the client program can query to obtain job status.

3.2.6 Data Services – SRB & Gridftp

There are two technologies employed primarily for data services within the system – the Storage Resource Broker (SRB - www.sdsc.edu/srb) and gridftp. SRB provides users with federated access to their datasets, this allows us to store, replicate and backup user's files with ease and to the user the files still appear as if in one place regardless of where they are physically stored.

The user stages their input files to the SRB resource pre-job submission and they keep a note of where they have saved the files. They then can incorporate the location of where they saved their job files in their job submission.

When ws-gram receives the user's job, it uses another service internal to GT4 called the "Reliable file transfer" (RFT) service to manage the staging in and out of files via the gridftp server. Natively Globus' Gridftp server does not support the SRB namespace so an additional gridftp-srb enabled server (http://www.globus.org/toolkit/docs/4.0/data/gridftp/GridFTP_SRB.html) has been deployed to enable this. Ws-gram stages the data and then submits the job to the scheduler adaptor. On notification of job completion, ws-gram once again uses RFT to stage out the files to SRB. The user can then download the files from their SRB space.

3.3 Other Services used by the workflow

The system uses multiple other "support" services to accomplish the workflow, for example an Authorisation, Authentication & Accounting service (AAA), delegation service, registry and various data services.

3.3.1 Authentication, Authorisation and Accounting (AAA)

The AAA service provides a single-sign on facility to the grid and authenticates the user and authorises the user to use various resources. The accounting mechanism tracks the user's quota. The delegation service provides the ability for the user to allow services that they trust to act on their behalf. GT4 provides the AAA mechanisms for the grid.

3.3.2 Registries

Registries are used for discovery of resources and services on the grid and for publishing job result sets. Other data services handle the transfer of data sets to and from computational resources and permanent storage of datasets. The system makes use of registries implemented using GT4's MDS for the dynamic discovery of resources and the ebXML registry (<http://www.ebxml.org/>) for both the publishing of data sets and discovery of services.

4 Example usage of the Architecture – pmd*CRC usage of the FastFlow service Workflow example

The Desktop Modelling Toolkit (DMT) is a software packaged developed under the pmd*CRC which orchestrates a demanding workflow used by computational geoscientists. The workflow requires integration of disparate data, computation and flexibility

for algorithm substitution, whilst providing an interface so that a non-expert user can get familiar with geological modelling.

Through the use of SEE and APAC Grid services and open standards for information exchange, the pmd*CRC's DMT has established a new workflow that aids in significantly improving linkage and the amount of research a investigator can conduct. This workflow is detailed in Figure 2.

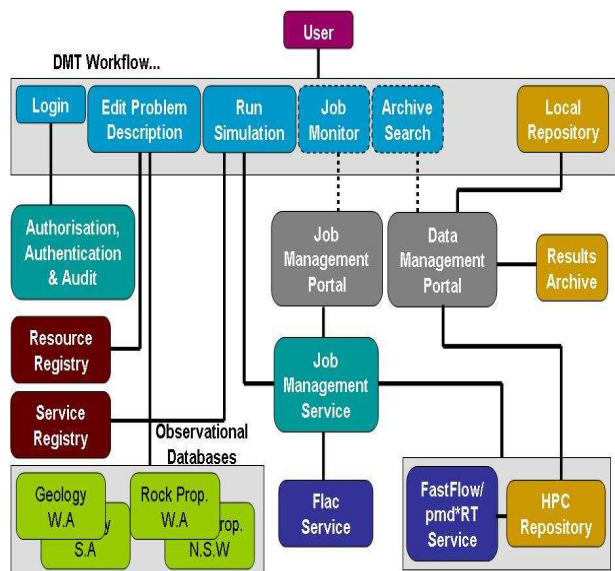


Figure 2 An example of the service interactions involved in the pmd*CRC modelling toolkit.

The DMT provides several user interfaces to aid a user in constructing a job to model a particular phenomenon. The toolkit integrates several other 3rd party tools that are used to create meshes and input parameter files. The toolkit also provides capabilities to convert these files between formats. A user uses the DMT to construct the job they require and once it is ready, through the interface they submit the job to the grid.

On submission, the toolkit uploads input data files to SRB on the user's behalf.

Then, in the instance a user is attempting to solve a problem which is best solved by a numerical algorithm available in Fast Flow, the job would be sent to the ws-FastFlow service on the grid. The client program creates a submission message which is ws-FastFlow compliant and sends it to the service. The service is listening for messages and once the message is received, processes the content – and sends a response in the form of an Endpoint Reference (EPR) to the job submitted (Gudgin and Hadley, 2005). The client can use this EPR to query the service for a number of things including – state of the job; retrieving the standard output & error; retrieving intermittent result sets and terminating the job.

ws-FastFlow service processes the job submission request, supplementing any “code-specific” criteria necessary such as site to run job at, queue & cpu requirements etc and finally constructs a ws-gram compliant job description document based on the job’s requirements. The Service then submits the job document to ws-gram in a ws-gram message. ws-gram processes the job, extracting the individual components. It firstly checks the job description file, stages input data and then submits the job to the scheduler adapter for conversion to a script the local scheduler can queue. ws-gram returns an EPR to the ws-FastFlow service as a reference to the ws-gram job.

The job gets queued by the local scheduler and executed whenever the first available resource that matches the job’s requirements comes available. Once the job has executed and completed, the local scheduler notifies ws-gram which then stages out the result files to SRB for the user. ws-gram notifies ws-FastFlow and when the Client next queries the Service (ws-FastFlow) with the EPR, they will discover that the job has completed and their result sets are available in SRB. The Client can now download and easily share the files from SRB when they desire.

5 Summary

Through the use of a service oriented architecture, the Geosciences support project in the APAC Grid project was able to achieve its goal of linking researchers in the community to computational and data resources. The project developed an architecture that not works in the Geosciences community, it could be used to grid-enable other communities. ws-gram based services were used to allow standards based communications between services in the grid environment. Researchers can now use services together to create simple workflows to solve extremely complex science problems.

6 References

Burbeck, S. (2000): The Tao of e-business Services: The evolution of Web applications into service-oriented components with Web services. Emerging Technologies, IBM Software Group. <http://www.ibm.com/developerworks/webservices/library/ws-tao> (accessed August 10, 2006)

Foster, I., Kesselman, C. and Tuecke, S. (2001): The Anatomy of the Grid. Intl. J. Supercomputer Applications

Foster, I., Kesselman, C., Nick, J.M. and Tuecke, S., (2002): The Physiology of the Grid, An Open Grid Services Architecture for Distributed Systems Integration. Argonne National Laboratory. Working

Draft: <http://www.globus.org/alliance/publications/papers/ogsa.pdf> (accessed August 10, 2006)

Gudgin, M. and Hadley, M., (2005): Web Services Addressing 1.0 – Core: W3C Working Draft 15 February 2005. WC3. <http://www.w3.org/TR/2005/WD-ws-addr-core-20050215/> (Access August 11, 2006)

Lin, W.A. (2005): Building a unified grid, Part 1: Grid architecture in the Telescience Project. National Center for Microscopy and Imaging Research. <http://www-128.ibm.com/developerworks/grid/library/gr-unified1/> (accessed August 10, 2006)

Srinivasan, L. and Treadwell, J., (2005): An Overview of Service-oriented Architecture, Web Services and Grid Computing. HP Software Global Business Unit http://devresource.hp.com/drc/technical_papers/grid_soa/index.jsp (accessed August 11, 2006)

Woodcock, R. and Fraser, R. (2005): APAC Geosciences Project: Linking Academia, Industry and Government to HPC. From SEE Grid II Conference, 2005, Canberra

Woodcock, R. and Fraser, R. (2005b): Towards Service-Oriented Geoscience: SEE Grid and APAC Grid. In Proceedings of Australian Academy of Science Elizabeth and Frederick White Conference: Mastering the Data Explosion in the Earth and Environmental Sciences, 2005, Canberra