The Integrated Marine Observing System: Delivering Reliable and Usable Services on the National Infrastructure

Dr Peter Blain 1, Dr Roger Proctor 2

1 University or Tasmania, Hobart, Australia, peter.blain@utas.edu.au
2 University or Tasmania, Hobart, Australia, roger.proctor@utas.edu.au

INTRODUCTION

The Integrated Marine Observing System (IMOS) is an Australian Government infrastructure project. IMOS is an integrated national array of observing equipment that monitors the open oceans and coastal marine environment around Australia, covering physical, chemical and biological variables. All IMOS data is freely and openly available through the IMOS Ocean Portal for the benefit of Australian marine and climate science. IMOS observations are guided by societal needs for improved ocean information, and focused through science planning undertaken collaboratively across the Australian marine and climate science community.

The IMOS Ocean Portal is an open source project first released in 2009, and most recently updated in March 2014. The latest redevelopment was motivated by two factors: Firstly, the user interface was difficult to understand. Secondly, the unreliability of the infrastructure and services that back the application were exposed to users. These were, and remain, difficult issues to manage given the complexity of the portal's requirements and the array of services that underlie it.

The problem, however, became more tractable following a migration to the National Servers Program (NSP) and the availability of fast storage. The redevelopment project attempted to extract maximum value from the new, albeit incomplete, infrastructure by placing an increased focus on methodologies and practices such as scrum, test management, usability testing, continuous integration and configuration management. The emphasis on such practices led to an improved product, both in terms of quality and usability.

SYSTEM OVERVIEW

The IMOS Portal web application is built on Grails and ExtJS. Grails is an opensource framework for web development using Groovy. Grails allows rapid development of web applications using approaches such as "convention over configuration". This approach allows Grails (rather than the developer) to manage the configuration of the underlying components. Grails is built on Java technologies such as Spring (for controller logic, run-time configuration, transactions etc) and Hibernate (for object-relational mapping) - all of which are managed by the Grails framework. Groovy (the programming language used by Grails) works with existing Java libraries. Grails uses the MVC (Model, View, Controller) design pattern, and encourages the separation of business logic into service classes. The portal's administrator interface is built with straight Grails. The client-side user interface is built using ExtJS, which is an opensource JavaScript library. The map functionality is implemented using OpenLayers, which is an opensource JavaScript library for displaying interactive maps.

The portal application is backed by a PostgreSQL database, and uses the Grails Object Relational Mapping (GORM) implementation to persist the domain objects. The search functionality, however, is delivered by GeoNetwork, which is a Java based metadata catalog for geospatial data. The IMOS information infrastructure group (eMII) is an active member of the GeoNetwork opensource community. The spatial search aspect of the portal's broader search capability has been further enhanced by a custom (Grails based) spatial search index that sits in-between the portal and the catalog. This was developed to ensure that search results from spatial queries not only overlap the bounding box supplied by the user, but that spatial features actually exist in the requested region.

Map layers for visualisation are delivered via servers supporting the Open Geospatial Consortium (OGC) WMS standard. The metadata records that describe the IMOS data conform to the ISO19115 Marine Community Profile. Non-gridded IMOS data is stored as NetCDF files and harvested into a postgres database. The non-gridded data is then served by Geoserver, which is an open-source OGC compliant WFS implementation. Because of its size, gridded data is retained in NetCDF format and served through proprietary web services.

USABILITY IMPROVEMENTS

The first project objective was to improve the usability of the portal's user interface. A three step wizard-like process was implemented that provides an easy to follow workflow. In step one, the user searches for data collections using a faceted search interface. Users are familiar with faceted searching, which is commonly used on e-commerce sites for hotel booking, car sales, retail sales, etc. Step two allows the user to visually subset one or more data collections. The visual nature of the map interface makes it easy for users to create a subset. In step three, the user can download the subsets in a choice of formats including CSV and NetCDF. The download step is a one-click operation.

The simple three-step workflow was intended to be intuitive, but this was not left to chance. Tests were conducted to verify the system's usability. Users who had not previously seen the new incarnation of the portal were asked to interact with the interface without prompting or guidance by the test moderator. These tests uncovered further usability issues, which were addressed and retested with fresh users.

The development work was carried out under the SCRUM methodology. SCRUM is an agile process that improves the rate at which user feedback can be fed into the development cycle. This is intended to improve engagement with user community so that usability issues can be rapidly corrected.

RELIABILITY IMPROVEMENTS

The second objective was to improve the reliability of the portal. A variety of techniques and practices were used to help improve the stability of the user experience.

An increased emphasis was placed on testing. Unit testing and integration testing of the server-side code was implemented using the Grails unit testing framework (which is built on the popular Java unit testing framework jUnit). The front-end JavaScript is tested using Jasmin. Functional and regression testing is performed every two weeks as part of the SCUM cycle.

Continuous integration (CI) is configured with Jenkins. CI allows code to be merged multiple times per day and helps reduce integration issues. It also forces all automated tests to be run repeatedly throughout the day.

Configuration management is implemented with Chef. This improves the management of IMOS servers within the NSP cloud environment. It also allows the infrastructure to be scaled through DNS load-balancing. This improves system performance and provides fail-over to parallel instances in the event of an outage. Parallel instances can run within the NSP or, potentially, in other data centres. IMOS will, for example, run instances of the core infrastructure at the Tasmanian NeCTAR cloud node. Even a complete outage at the NSP will be imperceptible to users.

All IMOS systems are monitored by Nagios. Monitoring and auto-recovery are intended to further enhance system up-time.

CONCLUSION

The IMOS portal is now easier to navigate, data can be downloaded more conveniently, and the web application is more robust. The sense of reliability and stability perceived by the user continues to improve as the national infrastructure stabilises and as the IMOS system is hardened through redundancy, caching, and other techniques that insulate users from underlying system issues.

INFORMATION ABOUT THE AUTHORS

Peter Blain is the information systems architect at IMOS. He leads a software development team that builds open-source software for the geo-spatial community. Peter's team is also responsible for building and maintaining the information infrastructure behind the Australian Ocean Data Network (AODN). Peter was previously the software development manager at the Tasmanian Partnership for Advanced Computing (TPAC). Prior to his move into e-research, Peter was a freelance analyst/programmer in the banking sector – both locally and internationally. Peter has a PhD in cognitive science, a Masters degree in accounting, and a Bachelor of Engineering degree with a major in computer systems.

Roger Proctor has been the Director of the eMarine Information Infrastructure facility of the NCRIS capability Integrated Marine Observing System since 2008. He is responsible for developing and managing the information infrastructure for discovery and access to IMOS data, and developing the wider Australian Ocean Data Network. He is on the Steering Committee of the EU-US-Australian Ocean Data Interoperability Platform project, co-chair of the Research Data Alliance Marine Data Harmonisation Interest Group, and is a member of the US-IOOS Data Management and Communications Steering Team. He has a degree in mathematics and a phd in numerical modelling. Prior to moving to Australia in 2008 he led the UK Natural Environment Research Council marine science programme 'modelling and observing systems for coastal seas'.