DMAS Critical Design Review - March 15, 2007

Background

Purpose and Scope of the Review

On March 15, 2007 NEPTUNE Canada and VENUS (the Projects) as a normal part of their risk management and quality assurance efforts invited a number of individuals to carry out a critical design review (CDR) of the proposed Data Management and Acquisition System (DMAS). This CDR meeting was a follow up to the previous Preliminary Design Review meeting held in Victoria June 2006.

The purpose of this review is to provide the DMAS team with an external evaluation of the work at a time when it is ready to embark on the implementation of a large body of code and services in support of the NEPTUNE and VENUS ocean observatories.

Mandate of the Review Team

The mandate of the review team therefore has been interpreted to look at the work progress, to assess the validity of the proposed architecture to support both the observatories and their stakeholders, to consider the current design choices and to rate the chances of success given the time line for implementation, the budget allocation as well as the suitability of the staff complement. In particular, the review team was asked that the following questions be addressed:

> Can the design, as disclosed, satisfy the System Requirements?
> Has the system baseline been established and documented to enable detailed design to proceed with proper configuration management?
> Are adequate processes and metrics in place (QA) for the program to succeed?
> Are the risks known and manageable for development and operational testing?
> Assuming that the resources would be adequate, is the program schedule realistic (technical risks)?
> Has the DMAS team adequately communicated with the scientific representatives of the Projects as well as with scientists outside the Projects?
> Is the architecture of the system proposed sound, forward-looking and flexible enough to serve the needs of the Projects for many years?
> Is the proposed overall system design adequate to support the architecture?
> Are the science use cases that are driving the design and architecture sufficiently representative of the needs of ocean scientists?
> Do you believe that DMAS, in its planned shape and form, will be capable of addressing the requirements of its stakeholders: academia, government, industry, and international users, for instance in the area of (meta) data standards?
> Can you recommend one or several meta data (or data description) standard(s) that DMAS should adopt to publish its data holdings?

Overall Assessment

Overall the CDR committee was very impressed with the progress made since the first PDR meeting with the implementation of many of the recommendations made from that meeting.

Major concerns with regard to QA have been addressed and the risks understood, although by the nature of scientific endeavour, they cannot help but be somewhat ambiguous and not well defined. A new science executive has been hired with the mandate of continuing to gather requirements from stakeholders. In general the committee feels the project is on track, and the design will meet the specified system requirements. We believe the program schedule may be a little optimistic and should be watched carefully, especially during the integration phase. The committee also applauds the choice of architecture with the move to web services and workflow, but has specific issues that should be further investigated as outlined below.

The use of a service oriented architecture (SOA) we believe is important because the needs and functional requirements of the ocean scientists are not well understood at this time. This is not a fault of the DMAS teams but reflects the reality of trying to define the functional aspects of any new project, whose capabilities are still unknown to the larger scientific community. SOA provides the greatest flexibility to change parameters and process when NEPTUNE and VENUS are in production, in light of changing realities and demands of the scientific community.

However, SOA does have a down side. In traditional software development processes, workflows and interactions between processes are largely hard coded and can be thoroughly QC tested before deployment. As result most of the development work and testing occurs before integration and deployment. With SOA, because of workflow and loose coupling of processes the logic and interaction between processes can not be tested until the integration phase or even later. That is why the committee expressed concerns about the integration timetable.

Committee Answers to specific mandate

Can the design, as disclosed, satisfy the System Requirements?
We believe at a conceptual level the system design is good. However greater comfort would be given to the committee and scientific users if detailed logical and physical design documents, UML diagrams, hardware diagrams, GANTT charts and other such information were made available for external review.

Are the technological choices made by the DMAS group sound and providing a way forward for NEPTUNE and VENUS in view of the challenge they represent?
The software choices appear sound. No discussion of the hardware choices were made. The committee thinks the DMAS team underestimates the work involved and time required in designing, selecting and procuring the hardware and integrating the components and the software, and this may have an impact on scheduling.

Are adequate processes and metrics in place (QA) for the program to succeed?
The committee believes significant progress was made in this area. The DMAS team are doing more QA than would normally be expected for a research project. However we believe the QA processes should be extended through the integration and operational phases and early testing with the research community. An example of early operation QA is the Large Hadron Collider software testing in CERN, for example, which went through a number of "service challenges" as it scaled from initial deployment to full production. For example the demand on the shore station for 100 simultaneous users did not seem to be clearly addressed. NEPTUNE might benefit by installing a beta-system at 5-10% of the real system and then gradually building the system up to full capacity. This would help with integration issues and provide experience to the operational team.

Are the risks known and manageable for development and operational testing?

The DMAS team seems to understand the risks. The biggest risk is the lack of detailed scientific functional requirements. But the choice of a SOA architecture should mitigate against that risk. The deployment of the DMAS on VENUS will provide early feedback on their assumptions.

Assuming that the resources would be adequate, is the program schedule realistic (technical risks)?

In general the program is realistic, but may be too aggressive with respect to the integration phase. This will have to be watched carefully. For example, there was no schedule presented for the installation of hardware at the two locations. Some of the DMAS team may need to be relocated to Port Alberni for a few months; this was not mentioned. Details of project timelines with key milestones were not presented to the CDR team making evaluation challenging.

Has the DMAS team adequately communicated with the scientific representatives of the Projects as well as with scientists outside the Projects for the purpose of collecting requirements?

The DMAS team has made significant steps to involve the researchers. However, the question of ownership of an instrument and data seems to be an unresolved issue. It is essential that the team maintain a strong link with the researchers as their requirements may evolve. The use of Wiki, blogs and other tools to communicate as much details of the NEPTUNE/VENUS architecture to the scientific community would be helpful.

Do you believe that DMAS, in its planned shape and form, will be capable of addressing the requirements of its stakeholders: academia, government, industry, and international users, for instance in the area of (meta) data standards?

The system seems to meet the needs of VENUS. NEPTUNE will be approximately 10 times larger than VENUS (a large but not huge increase). The use of the NRC/HIA system has addressed the needs of the astronomical community. The data used by the ocean community is very similar in structure so the system should work.

General Recommendations

The following general recommendations are those that were identified by the committee at the March 15 meeting. Subsequent to that meeting additional specific recommendations have been added by the individual members of the committee at the end of this report and should also be addressed by NEPTUNE/VENUS management.

Overall the committee was very pleased with the progress of the DMAS and the following recommendations should not be seen as being critical of the project, but instead as suggested ways of fine tuning the management and monitoring the project.

1. Carefully review the integration timetable and prepare budgets and schedules to allow for the addition of resources and expertise as required for specific project activities such as logical and physical design, integration testingand system deployment . One of the committee's main concerns was the tight timeframe for the integration phase of the project, as such this should remain an important focus area for management.

2. Extend QA processes through integration and production. The committee was impressed by the appointment of an external QC person, but recommends their responsibility be

extended to the integration and early operational phase especially in terms of QC documentation and processes. SOA moves many software QC design issues from the coding phase of the project to the integration and sometimes operational phase.

3. Make the design and development Wiki and other tools open to the scientific community or the public. The committee was disappointed that no detailed architecture documents, UML diagrams, and so forth were not tabled at the CDR. An architecture and methodology needs to be demonstrated to deal to with poorly defined functional requirements. By employing a largely process and metadata driven solution this can help mitigate the nebulous needs of the scientific community and changes in ocean science technology. Of particular note was the absence of the following key architectural UML work products: Class and/or Object Diagram, Component Diagram, and Activity Diagram.

4. Ownership of instruments and data still seems to be a sensitive issue. Concern remains with respect to the integration of the observatory with existing realtime networks. While the current design has more clearly addressed this issue at a relatively high level, the actual manner in which this will be achieved continues to be unclear. Of major concern is the control of the instruments and how that will be accomplished. From the perspective of these networks, the influence of NEPTUNE on an instrument should be completely transparent. The presentation suggested that this would be the case, but not convincingly so. The committee recommends that the DMAS teams prepare a couple of detailed scenarios of how instruments will be deployed and accessed by researchers and how data will flow unimpeded to the research community. Two suggested scenarios are the installation, control and management of a flow cytometer, and secondly the data flow of seismic sensor array to the NRCan Pacific Geology Centre.

5. For security management the committee recommends use of proven directory/certificate services rather than user name/password for security. A centralized directory for security services can be leveraged by other IT systems. Providing a single source for user credentialing (i.e. Authentication, Authorization and Access) can in turn be consumed by the RDBMS, web application servers, unstructured data/file storage, the backup, archival and restore infrastructure will greatly reduce the administrative overhead, access controls as well as the revocation of privileges.

6. The reliability targets are not realistic. 99.99% for a research project seems very optimistic. A more realistic advertised target should be 98-99%. These numbers do affect the technical design and should allow DMAS more flexibility in the hardware and software design and also, if necessary, de-scope components of the project. The committee recommends that the DMAS start with less aggressive targets on all aspects of the projects and improve as demand and resources warrant. If necessary the committee recommends increased investment and pursue high availability at the shore station with more modest investment in the data centre at least initially.

7. The committee expressed some concerns about the use of the NRC/HIC database. It is recommended that the DMAS team carefully investigate any issues with respect to lack of referential integrity mechanisms between structured and unstructured data in the proposed database. They should also insure that mechanisms are being created in the data model to ensure referential integrity and security of file based information (no direct access is allowed to the file system except through the web services/database). Also they

should make sure that there are tools available for doing index rebuilds in the event of logical corruption of the indexes.

8. The DMAS management team needs to clarify the role of UVic computing in the Project and whether the University is expecting NEPTUNE to put its resources into the new UVic computing room. The CDR committee can understand that DMAS wish to run their own computing systems but there is a large computing centre very close with lots of expertise. The UVic central computing group are also knowledgeable about grids and databases. NEPTUNE management may consider retaining innovative activities such as software development but outsourcing managed operations of IT systems to an external party (e.g. University of Victoria). An example of this would be NEPTUNE employing a Subject Matter Expert for level 1 help desk support (which requires domain expertise) while outsourcing level 2/3 help desk support (information technology skills). A cost/benefit analysis of the two models may be beneficial in helping to arrive at the best and most cost effective support model for NEPTUNE.

9. The committee suggests that the DMAS team combine the overall Spiral development methodology with Model Driven Architecture (MDA) and Unified Modeling Language (UML). MDA defines an approach to architecture that separates the specification of system functionality from the implementation of that functionality on a specific technology platform. To this end, the MDA defines an architecture for models that provides a set of guidelines for structuring specifications expressed as models. The MDA approach and the standards that support it allow the same model that specifies system behaviour to be realized on multiple platforms through auxiliary mapping standards, or through point mappings to specific platforms. This allows different applications to be integrated by relating their underlying models, which in turn enables integration and interoperability as well as system agility as technologies inevitably evolve over time.

10. The committee recommends that the DMAS team do further investigation and research into various workflow tools such as BPEL and use of mainstream WS standards. Kepler is a useful for the scientific community, but we are concerned about its ongoing evolution and whether it will fit within a broader SOA environment. Modeling tools and infrastructure that support BPEL for experimental or research work flows are starting to be available. The UK eScience program and the Globus team are both developing BPEL workflow tools for the scientific community. BPEL provides an XML notation and semantics for specifying process behavior based on Web Services. BPEL orchestrates Web Services by specifying the order in which it is meaningful to call a collection of services, and assigns responsibilities for each of the services to partners. WS-Notification is a family of specifications that standardize the process of creating an event-based system in a Web services environment. They were developed in conjunction with the Web Services Resource Framework (WSRF), which provides a way to use state in a stateless environment, but they can also be utilized by plain Web services. WSN consists of three specifications; WS-BaseNotification, WS-Topics and WS-BrokeredNotification.

11. Some concerns were expressed about the data rates and storage. The data rates are modest (60 GB per day) and up to 3 TB/day for HDTV; with and expected order of 100 TB/day compressed. Rather than compressing raw data it can be stored on tape. NEPTUNE maybe should adopt a simple backup system for the data (copy the data to tape and store it offsite). UVic has a contract with a company for such a service.

12. The DMAS's team proposed use of the grid is nice but may not be a high priority until the other elements of the system are operational. This is a project that might be identified as something to be de-scoped if there are budget or schedule overruns. Instead it would be prudent for them to have a small cluster for dedicated NEPTUNE processing. They should be aware that the grid is more challenging if commercial (licensed) software is used.

13. The committee suggest that the DMAS team reconsider the presence of a mediation layer in the overall solution architecture. Lack of an ESB could introduce additional effort rather than what can be acquired at little to no costs:
    a. Management of pub/sub topic at the bus level
    b. Rudimentary business rules and filters instantiated at the bus level
    c. Reduction in brittle point to point integration rather than mediating these exchanges through compiled code (asynchronous as well as synchronous)
    d. Potential increase risk of new technology adoption, additional training and some reported performance issues in adopting the ESB (touch base with WESB-HL7 IBM team for insight on this issue)

    Some of the performance concerns of the Enterprise Services Bus (ESB) can be mitigated as well as the risk of creating functionality that it inherent in adopting an ESB. Further employing a Service Registry and Repository will help resist point-to-point integration, introduce change controls, amplify the opportunities of collaboration in the larger research community and improve the overall robustness of the solution

The benefits that would accrue to the NEPTUNE project in the near term adoption of an Enterprise Service Bus would include:

Reduction of point-to-point interconnections and simplification of application programming through separation of integration logic from the applications and from process logic
A publish-and-subscribe matching engine configurable at run-time rather than at compile time which is capable of routing information in real time based on topic and content to any endpoint
Message validation and transformation amongst possible different message formats, including Web Services, other XML and non-XML formats
Ability to route messages based on domain specific (research or experimental) rules in order to action on events or to match information content, processes or for correlation via Complex Event Processing (CEP)
Mediates (provides routing, transformation and logging) between like or dissimilar service requesters and providers (i.e. Web Services Definition Language-WSDL, Simple Object Access Protocol- SOAP, Java Message Service-JMS, HTTP-based Representational State Transfer-REST)
Publish/Subscribe topic instantiation and management
Simplified systems integration, ongoing development and reduction in the overall volume of source code to maintain

There were some initial concerns of the ability for the ESB to meet the non-functional requirements of data acquisition. Based on some basic configuration changes (i.e. the persistent subscriber ran in a transaction with the same commit count as the publisher) to a candidate ESB asset (i.e. IBM Advanced ESB) the following results were obtained:

Configuration	Rate (msg/s)
Persistent Pub/Sub with transaction	74
Persistent Pub/Sub w/ transaction, commit count=2, subscriber w/ transaction, commit count =2	141
Persistent Pub/Sub w/ transaction, commit count=3, subscriber w/ transaction, commit count =3	192
Persistent Pub/Sub w/ transaction, commit count=10, subscriber w/ transaction, commit count =10	425

These revised performance results demonstrated that a optimized mediation layer can effectively deliver and surpass the transaction throughput required by critical scalar sensor values.

Additional observations include:

The more messages included in a transaction the better the throughput.
FASTPATH bindings have a noticeable impact on non-persistent performance (50% increase)
SingleWrite increases the throughput of persistent messages by as much as 50% rather than the default WebSphere MQ setting of TripleWrite (results in a reduction of disk I/O)
A higher throughput of persistent publications can be achieved if the one persistent subscriber is also getting messages within a transaction and with a similar commit count to the publisher
A significant impact to overall ESB performance is the disk sub-system. High performance hard drives with modest RAID controllers (large cache with battery backup) and Fibre Channel interfaces will greatly improve message throughput and reliability

14. One of the project risks identified is staff turnover and advice was requested from the CDR team for mitigation strategies. With a small development team, the risk of losing knowledge and impacting timelines is considerable. Victoria is a small market for software development and market conditions are currently very favorable for employees. Here are some points to consider:

    Ensure software development processes and system design documents are well documented and kept current so that new resources may be ramped up quickly. Identify who is responsible for ensuring documentation be kept up to date.
    Adopt new technologies such as ESB to create and maintain a stimulating environment for software developers.
    Ensure that salaries are competitive with current market conditions. If budgets are not sufficient for base salaries to be competitive then make sure that there are other perks available to staff such as flexible work hours, dedicated training time, opportunities for aquiring new skills etc.

CDR Committee Members

Bill St. Arnaud, CANARIE (chair)
Paul Ross, BC Provincial Health Services Authority
Richard Baldwin, NRCan
Randy Sobie, University of Victoria, Department of Physics
Rob Heuchert, IBM Canada