A Reference Model for a Water Resources Observation Network

O'Hagan, R.G.[1]. R. Atkinson[2], S. Cox[3], P. Fitch [1], D. Lemon [1], and G. Walker[4]

[1] CSIRO Land and Water, CSIRO, Canberra
[2] Social Change Online Pty. Ltd., Sydney
[3] CSIRO Exploration and Mining, CSIRO, Perth
[4] CSIRO Information and Communication Technology, CSIRO, Canberra
Email: Rochelle.OHagan@csiro.au

Keywords: *Spatial data infrastructure, water information, distributed architectures, interoperability*

EXTENDED ABSTRACT

Meeting Australia's current and future water challenges requires timely access to the current and forecast future status of water resources. This in turn requires access to data, processing algorithms, hydrological models and resulting information products used by several hundred agencies in the water resources sector.

The Water Resources Observation Network Reference Model describes a process for developing a distributed interoperable framework for sharing information. This framework will enable the linking of data and processing services to forecasting and reporting technologies which will improve the visibility, currency and usability of water resources information. The Water Resources Observation Network (WRON) will allow for improved environmental, social and economic outcomes with respect to water.

Development of the WRON Reference Model is guided by a set of policies, and defines several architectural principles to ensure the feasibility, flexibility, extensibility of the WRON. The policies include 'Adopt, Adapt, Invent', recognising the benefit of working, wherever possible, with existing standards, protocols or procedures, extending or adapting them if necessary, and only inventing where adoption or adaptation is not possible. A second guiding policy is that the WRON Reference Model should not contain anything that cannot be implemented – thus ensuring the practicality and usefulness of the framework. Architectural principles include 'adequate description' to ensure that sufficient information is available to use a resource, 'subscribe, not describe' – to encourage reuse of published service descriptions and minimise effort in establishing new services, and 'no private contracts' which prohibits unpublished or private agreements between components that would compromise the interoperability of the WRON.

This paper introduces Version 0.1 of the WRON Reference Model and further describes the principles and policies outlined above. It also details the key components and information artefacts that compose the Reference Model, including registries to enable discovery, services and their descriptions or profiles, authoritative copies of data or caches and domain models for defining concepts within the domain. The paper also discusses the role of standards within the WRON and the need for governance.

1. INTRODUCTION

Meeting Australia's current and future water challenges requires excellence in water management. This, in turn, will rely on the availability of key information including:

- A synoptic understanding of the current status of water resources; and
- A reliable forecast of the probable future status of water availability.

Determining the current and future status of water resources, necessitates access to data, both current and historical, from a wide range of sources. These include in-situ and remote sensors, users, management authorities, government agencies and water companies. Additionally, access to hydrological models is needed in order to analyse the data and produce water information products.

The Water Resources Observation Network (WRON) will enable the coupling of data and processing services to a new generation of forecasting and reporting technologies greatly improving the visibility, currency and usability of information on Australia's water resources.

The WRON Reference Model (WRON-RM) describes an interoperability framework or 'spatial data infrastructure' for:

- linking Australia's many water and water related data assets;
- harnessing new data streams from satellites and on-ground sensor networks; and
- processing and utilising water information.

The framework is designed as a set of components brought together as a loosely coupled system based on distributed services conforming to standard interfaces.

In this paper, we describe key elements and artefacts of the WRON Reference Model. Section 2 provides background on the need for the WRON and how it will be used. Section 3 introduces the policies and architectural principles that will underpin the implementation of the WRON. In Section 4, the content of the Reference Model is described in more detail. Section 5 discusses implementation progress to date and activities being undertaken for Reference Model Version 1.0. Finally, Section 6 provides conclusions.

2. BACKGROUND

The Australian water information community consists of several hundred agencies who collect and maintain water data. In addition, other users, both government and private, require access to data for forecasting, prediction and water management.

Current systems for information sharing rely on individually handled requests for supply of water data, using transfer technologies such as exchange of physical media (for example CD-ROMs), email, FTP or on-line delivery of data tables. Consistency in data formats is very low, and metadata often insufficient or even non-existent. Semantic interoperability in such an environment is very difficult to achieve without significant and ongoing transformation of data.

2.1. Why a WRON SDI?

Spatial Data Infrastructures (SDIs) are interoperability frameworks designed for spatial data. A SDI provides "a basis for spatial data discovery, evaluation, and application for users and providers within all levels of government, the commercial sector, the non-profit sector, academia and by citizens in general" (Nebert, 2004). SDIs facilitate interoperable access to data by encouraging data harmonisation and the adoption and use of standards.

The WRON Reference Model describes a spatial data infrastructure for water information. Based on internationally recognised standards, the WRON will allow discovery of water information resources and access to data and processing services. A common domain model will encourage harmonisation of data and reduce the cost of data sharing by removing tedious and time-consuming data transformation. Explicit governance arrangements for various elements, often delegated, play a key role in both system development and maintenance.

3. POLICIES AND PRINCIPLES

The WRON-RM employs several policies which guide its development, and defines architectural principles for the WRON framework. These policies and principles are in place to ensure the feasibility, flexibility and extensibility of the WRON.

3.1. Adopt, Adapt, Invent

Recognising the value of previous work and the advantages of working with recognised standards,

the WRON-RM adopts the principle that where possible, existing standards, protocols and procedures should be adopted for use by the WRON. If wholesale adoption is not possible, attempts should be made to adapt existing standards, protocols and procedures. The development of new standards, protocols and procedures should only occur when adoption and adaptation are not possible, in particular where WRON clearly has a new or unique mandate. Within the WRON-RM, this policy is referred to as 'Adopt, Adapt, Invent'.

3.2. Nothing that cannot be implemented

The 'Nothing that cannot be implemented' policy states that the WRON-RM will only describe in detail components that can be implemented and for which a reference implementation exists. The WRON-RM may identify but will not describe aspirational components.

3.3. Adequate Description

The first architectural principle defined by the WRON-RM is the concept of 'adequate description'. In order for data and/or processing resources to be usefully employed, the user must be able to gain sufficient knowledge about the resource to understand its use. For example, a query that requires the end user to nominate a time period as a calendar month must fully specify both acceptable formats and range. This becomes increasingly important when automatic service composition is considered – services need to be adequately described in order for semantic reasoning to be applied and the appropriate service selected and executed.

3.4. Subscribe not describe

There is a burden of effort associated with the need to describe all artefacts within a distributed environment. This burden is encountered in several areas including:

- Discovering many descriptions of individual resources and comparing them to understand the differences and potential suitability for a purpose;
- Difficulty and overhead of providing adequate description when recreating all aspects of metadata without re-use; and
- Difficulty of creating software components able to interpret and consume services based on ad-hoc descriptions.

Essentially, this means that the effort of description, while manageable for one artefact, becomes excessive when the system is scaled to hundreds or thousands of artefacts. In order to ease the burden of adoption, the concept of 'Subscribe, not describe' is introduced.

The principle of 'Subscribe, not describe' encourages service providers to implement, where possible, services that adhere to published descriptions rather than deploying a service that requires a new description. Where subscription is not possible, service providers are encouraged to publish their description so that others may subscribe to it.

This principle recognises that it is far more efficient to add a new resource by attaching it to a reusable description (for example a service profile) than to describe it fully.

3.5. No Private Contracts

By basing itself on published standards, the WRON-RM ensures that all WRON components will be discoverable, usable and interoperable. The existence of unpublished, private agreements between components that contradict or compromise published standards threatens the interoperable nature of the WRON. Thus the WRON-RM declares the principle of 'No Private Contracts'.

4. THE WRON REFERENCE MODEL

The WRON Reference Model (WRON-RM) describes the architecture for the WRON, key information artefacts, governance regimes and procedures, a set of system use cases and the standards that have been adopted for the WRON.

The WRON-RM follows the structure of the Reference Model for Open Distributed Processing (RM-ODP) developed by ISO and ITU-T (ISO/IEC, 1994). The RM-ODP provides a framework for the standardisation of open distributed processing and defines a specification for open distributed systems consisting of five viewpoints – Enterprise, Information, Computational, Engineering and Technology.

In the WRON Reference Model, the WRON architecture is described according to these viewpoints. In addition to the WRON Reference Model document, a UML model of the WRON architecture has been developed. This model is maintained in parallel with developments to the WRON-RM document and serves as an easily

distributed, formalised representation of the reference model.

4.1. Key Information and Computational Artefacts

The WRON-RM describes an architecture for the WRON itself. The architecture is based on a Service-Oriented Architecture (MacKenzie *et. al.* 2006) with web services providing functionality for data access, processing and service chain composition.

This architecture recognises that the information and the services that expose this information to the WRON will be owned, managed and contributed by a range of organisations.

Key components of this architecture are:

- Registries
- Services
    - Data
    - Processing
    - Orchestration
- Caches (with synchronization/ update mechanism)
- Service Profiles
- Domain Models

Registries

An important interaction for many stakeholders in the WRON is discovery. This includes discovery of:

- services that meet specified requirements; and
- definitions.

The need to discover implies the need for registries and registers (Hasselmeyer 2005), key artefacts of the WRON.

Each register requires a formal governance regime (c.f. ISO 19135 Geographic information – Procedures for item registration (ISO 2005)) determining who has rights of access for both reading and writing. Registries will need separate governance to establish individual registers and delegate governance to appropriate owners. This will also apply to semantic registries and their content including, for example, domain models, service instances and meta-information.

Services

Users will interact with the WRON through web services. They are likely to be exposed to a variety of service types including data services, processing services, registry services, orchestration services and metadata services.

The WRON-RM will define the requirements for exposing a service to the WRON.

Strong Forward Caches

Data Warehouses play a key role in spatial data infrastructures. They provide a performance optimization to expose data services. This requires the data that populates the warehouse to be available from stable sources, and the warehouse is kept up to date. In this context the warehouse can be viewed as a strong forward cache: an authoritative and readily accessible copy of the data.

An analogy of this is the credit card system. All credit card details for VISA holders are stored by VISA in the USA – at the 'point of truth'. An authoritative copy of this information is held by a bank in Australia or New Zealand and updated when necessary with, for example, changes to credit limits, or transaction information. These bank copies are forward caches, with a regular mechanism for synchronization and update relative to the point of truth.

Service Profiles

A service profile describes the minimal conformance level for a deployed instance of a service. This conformance may include:

- Service type (e.g. a data service);
- Content type exposed (e.g. flow data);
- Queries/invocation messages supported (e.g. 'get data for period');
- Quality of Service; and
- Documentation (metadata) requirements (e.g. data description, service metadata).

A service profile's primary role is to ensure that the expected behaviour of a service meets the expectations of the service consumer. It is also critical for:

- Documenting service behaviours in a reusable fashion;
- Providing service providers with an implementation checklist;
- Supporting automated conformance testing;

- Allowing system managers to assess compliance; and
- Creating similarity between different types of services to minimise overall system complexity (from both implementation and usage perspectives).

Service profiles in the WRON-RM are designed following and extending the concepts defined in ISO 19106 Geographic information – Profiles (ISO 2004).

Domain Models

A domain model provides authoritative definitions of concepts within the domain. When all services are consistently bound to the same domain model, semantic interoperability is achieved. That is, if all members of a community agree on the definition of a particular term or set of terms, then information using these terms can be exchanged between members of the community without the risk of misinterpretation.

Profiles, as containers of this binding, use concepts that are found within the domain model. That is, a service that subscribes to a particular profile that includes data must use data types defined within the domain model.

The domain model may comprise elements of other domain models and will need to reference these parts at their 'point of truth' or original location. As these referenced parts may be under governance by a responsible body (such as a standards body), there is a versioning requirement as the 'point of truth' may change in a way that invalidates the domain model. This requires orderly governance of the domain model.

A domain model is being developed for the WRON based on ISO standards and the Observations and Measurements pattern (Cox 2006).

4.2. Standards framework

In order to achieve data and service interoperability between the very large number of members of the water resources community, standardisation is key. A common data model and set of vocabularies provides the basis for standardised data description (metadata), and the development of standard definitions and descriptions of services ensures functional integration of services developed by different parties.

Core to these activities is the concept of 'WRON Compliance' which declares the goal of building the WRON from independent components under the management and control of various agencies, built on standard, open interfaces as defined by the WRON-RM. A compliance framework developed as part of the reference model implementation will provide a means for contributors to determine whether their services are WRON compliant, and hence interoperable with other WRON components. WRON compliance will be a pre-requisite for inclusion in WRON registries for service discovery and use.

The WRON will be built on top of suitable existing standards. These standards include generic cross-domain specifications (for example, web service interface protocols), and compatible existing subject-specific standards (for example, domain models) that satisfy particular WRON outcomes. Standards to be considered are those where evidence of implementation and capability within the sector can be identified, thus complying with the policy of 'Nothing that can't be implemented'.

Where suitable standards do not exist, WRON artefacts may be implemented without standards to guide all aspects of the implementation. This is expected to result in feedback to appropriate standards bodies.

The first set of standards to be considered is the ISO 19100 series. These establish a useful methodology for the formalisation and governance of a domain model (aka application schema) and encodings. This includes provision of a set of important component types for geometry, temporal objects, coordinate reference systems, and coverages. These 'horizontal' components are applicable across most geospatial information applications. Applications that share these components are interoperable at this level.

4.3. Governance

The WRON encompasses the whole community of organisations with an interest in water information and management in Australia. Governance policies are key in engaging this community and ensuring participation in the WRON. The WRON essentially implements an agreement amongst participating custodians and clients to use a particular set of standards and concepts. Each information type and each information artefact needs to have a transparent and effective governance regime and lifecycle.

In a practical sense, governance of the WRON means management and version control of information artefacts such as information and

domain models, controlled vocabularies and service definitions as well as the standard security concepts such as authentication, authorisation, accounting and audit (AAAA). Governance in the WRON is primarily realised using registries and registers of the various information artefacts.

Even the WRON Reference Model needs governance at both a policy and implementation level. At a policy level, this means that the Reference Model needs to be owned and maintained. At an implementation level, version control and configuration management is currently in place using a Subversion repository.

4.4. System Use Cases

The WRON-RM is scoped according to the need to deliver information products from multiple sources to an evolving set of business functions. It deals with all aspects of the implementation of such a system, including data delivery, use and long term management of a system that integrates a growing set of components.

A set of System Use Cases have been identified in the WRON-RM to describe the architecture according to the perspectives of different types of stakeholders and necessary management roles. This recognises that the various users of the WRON have different and possibly competing needs and outcomes that must be addressed. By describing the architecture from each perspective, the reference model ensures that each user's needs are met. The intention is that the WRON-RM should be accessible to a particular audience through a perspective that encapsulates as simply as possible the interaction that the audience will have with the WRON.

The set of perspectives that have been identified are:

- End User;
- Data Provision;
- Functionality Provision;
- Enablement and Governance;
- Cross-Business Domain Integration; and
- System Maintenance.

The WRON-RM ensures that each of these perspectives is consistent with the others and all infrastructure required to meet the expectations of these perspectives is described.

5. CURRENT WORK – TOWARDS VERSION 1.0

The WRON-RM team is currently working towards the release of Version 1.0 of the WRON-RM. A preliminary version of the reference model has been released to a limited audience, and the focus of current work is towards validating the concepts and methodologies described in WRON-RM 0.1 and developing reference implementations.

Key activities include defining a scenario mapping methodology, developing and deploying a registry, implementation of service profiles, development of a domain model and definition of governance policies.

6. CONCLUSION

This paper provides a brief overview of the Water Resources Observation Network Reference Model. It is recognised that standards play a vital role in ensuring the interoperability goals of the WRON are met, and this is addressed in the WRON-RM through architectural principles as well as in the adoption of ISO and OGC standards. Although the WRON-RM is in the early stages of development, the key concepts have been identified and progress is being made on reference implementations.

7. REFERENCES

Cox, S., (2006), Observations and Measurements, Open Geospatial Consortium Best Practices Document OGC 05-087r4.

ISO/IEC CD 10746-1 (1994), Basic Reference Model of Open Distributed Processing – Part 1: Overview and Guide to Use.

ISO/IEC DIS 10746-2 (1994), Basic Reference Model of Open Distributed Processing – Part 2: Descriptive Model.

ISO/IEC DIS 10746-3 (1994), Basic Reference Model of Open Distributed Processing – Part 3: Prescriptive Model.

ISO/IEC CD 10746-4 (1994), Basic Reference Model of Open Distributed Processing – Part 4: Architectural Semantics.

ISO 19106 (2004), Geographic information – Profiles.

ISO 19135 (2005), Geographic information – Procedures for item registration.

Hasselmeyer, P. (2005), On Service Discovery Process Types, in Proceedings of 3rd International Conference of Service Oriented Computing, Amsterdam, 2005.

MacKenzie, C.M, Laskey, K., McCabe, F., Brown, P.F., and Metz, R., (2006) OASIS Reference Model for Service Oriented Architecture V1.0, OASIS Open.

Nebert, D.D. (2004), Developing Spatial Data Infrastructures: The SDI Cookbook, *Global Spatial Data Infrastructure*.