NATIONAL RESEARCH FLAGSHIPS

Water for a Healthy Country

The Water Resources Observation Network

Reference Model - Version 0.1

Lemon D.
Cox S.
Walker G.
Atkinson R.
Fitch P.
O'Hagan R.

April 2007

CSIRO

Water for a Healthy Country

The Water Resources Observation Network

Reference Model – Version 0.1

Lemon D.
Cox S.
Walker G.
Atkinson R.
Fitch P.
O'Hagan R.

April 2007

Water for a Healthy Country is one of six National Research Flagships established by CSIRO in 2003 as part of the National Research Flagship Initiative. Flagships are partnerships of leading Australian scientists, research institutions, commercial companies and selected international partners. Their scale, long time-frames and clear focus on delivery and adoption of research outputs are designed to maximise their impact in key areas of economic and community need. Flagships address six major national challenges; health, energy, light metals, oceans, food and water.

The Water for a Healthy Country Flagship is a research partnership between CSIRO, state and Australian governments, private and public industry and other research providers. The Flagship aims to achieve a tenfold increase in the economic, social and environmental benefits from water by 2025.

The work contained in this report is collaboration between CSIRO Land and Water, Information and Communications Technology Centre and Exploration and Mining

© Commonwealth of Australia 2007 All rights reserved.
This work is copyright. Apart from any use as permitted under the Copyright Act 1968, no part may be reproduced by any process without prior written permission from the Commonwealth.

Citation: Lemon D, Cox S, Walker G, Atkinson R, Fitch P and O'Hagan R. 2007. Water Resources Observation Model – Reference Model Version 0.1. CSIRO: Water for a Healthy Country National Research Flagship. 93 pp

DISCLAIMER

CSIRO advises that the information contained in this publication comprises general statements based on scientific research. The reader is advised and needs to be aware that such information may be incomplete or unable to be used in any specific situation. No reliance or actions must therefore be made on that information without seeking prior expert professional, scientific and technical advice. To the extent permitted by law, CSIRO (including its employees and consultants) excludes all liability to any person for any consequences, including but not limited to all losses, damages, costs, expenses and any other compensation, arising directly or indirectly from using this publication (in part or in whole) and any information or material contained in it.

For more information about Water for a Healthy Country Flagship visit or the National Research Flagship Initiative at www.csiro.au.

Preface

In late 2005, a group of CSIRO researchers funded through the eWater CRC started work on the first official project associated with the Water Resources Observation Network (WRON) concept. At that time, whilst much discussion of the WRON and its goals had been held, there was little understanding of how it might function. It was the goal of this project team to describe an architecture, based on best practice at the time, to realise the vision of the WRON.

By early 2007, the project team, with some additional members, had come to a point where they believed they had gained enough insight into the requirements of the WRON to be able to produce a framework for a reference architecture or model. As such they wrote the following document calling it Version 0.1 in recognition that it was not complete and further work was required. The WRON – Reference Model (WRON-RM) 0.1 was delivered in April 2007.

On review of the content of the WRON-RM 0.1, the project team realised that considerable further work was going to be required to develop a complete Reference Model. This was due to two factors:

1. It was recognised that use cases described in the document did not fully reflect the WRON and that the true use cases for the WRON were not well understood. As such, considerable work was going to be necessary to discover, describe and analyse these use cases in order to understand their impact on the requirements of the WRON.
2. A number of key components (both informational and computational) described within WRON-RM 0.1 did not exist and would require more thought and experimentation before they could be operationalised. This meant that a key principle of the Reference Model 'Describe only what can be shown to be implementable' could not be met.

These discoveries did not deter the work. To the contrary they have led to the development of a broader research program under the Water Informatics Research Stream to realise the vision of a distributed, robust, evolvable water information repurposing system. However, they did result in the WRON-RM 0.1 document not being published in any form. This decision was agreed to by all project team members in recognition of the incomplete nature of the work.

In the intervening years, members of the original WRON-RM project team and others associated with the Water Informatics Research Stream have found themselves regularly returning to the WRON-RM 0.1 for guidance. In particular, many of the principles and high level requirements remain as valid today as they did four years ago. Furthermore, a number of groups, both national and international, have approached CSIRO seeking advice around the subjects discussed within the WRON-RM 0.1. As such, the original WRON-RM team have come to re-consider the original decision not to publish WRON-RM 0.1 and have now agreed that, even in its incomplete state, there are elements within in it worthy of publication.

With this in mind, readers should be aware that the following document was completed some four years ago and has not been updated since. The authors recognise that elements within the document are incomplete or, are now, invalid. In particular, the context within which the original document was written has changed considerably. Eastern Australia is no longer in the grip of drought and there is now a national agency responsible for the collation of water information and reporting on the state of the nation's water resources. However, the high level principles and requirements remain as valid today as they did when originally penned. As such this document provides a valuable contribution to the ongoing discussion about future water information systems.

David Lemon
Research Stream Leader – Water Informatics
March 2011

Table of Contents

Preface ..... **i**
Table of Contents ..... **ii**
Executive Summary ..... **iv**
1. Introduction **1**
1.1. Background 1
1.1.1. The WRON 2
1.2. The WRON Reference Model 3
1.2.1. Principles of the WRON-RM 3
1.2.2. Governance of the WRON-RM 3
1.2.3. Structure of the WRON-RM 4
1.2.4. Evolution of the WRON-RM 5
1.2.5. Scope of WRON-RM 0.1 6
1.3. Context of the WRON-RM 6
1.3.1. Spatial Data Infrastructures 6
1.3.2. GRID Computing 7
1.3.3. Simulation Frameworks 7
1.3.4. Water Data 8
1.4. The Dam Level Scenario 8
2. Notation **10**
2.1. Glossary 10
2.1.1. Abbreviations 10
2.1.2. Terms and Definitions 11
2.2. UML Diagrams 18
2.2.1. Use Cases 18
2.2.2. Data Models (Class Diagrams) 18
3. References **19**
3.1. ISO Standards (or Drafts) 19
3.2. OGC Best Practice Documents 19
3.3. Documents and Books 19
4. Enterprise Viewpoint **20**
4.1. Overview 20
4.2. Business Requirements for the WRON 20
4.2.1. Risk Mitigation 21
4.2.2. Contribution to the WRON's Objectives 21
4.2.3. Alignment with Other Initiatives 23
4.3. Architectural Requirements and Principles for the WRON 26
4.3.1. Implementation Neutrality 27
4.3.2. Timeliness, Transparency, Traceability 27
4.3.3. Scalability 28
4.3.4. Simplicity 29
4.3.5. Adequate Description 29
4.3.6. Who pays the cost? 30
4.3.7. Semantic Robustness 31
4.3.8. No Private Contracts 31
4.4. Governance 31
4.4.1. Key Governance Regimes 31
4.5. WRON System Use Cases 33
4.5.1. End User Perspective 34
4.5.2. Data Provision Perspective 40
4.5.3. Functionality Provision Perspective 44
4.5.4. Enablement and Governance Perspective 50
4.5.5. Cross-Business Domain Integration Perspective 50
4.5.6. System Maintenance Perspective 50
5. Information Viewpoint – Key Artefacts **51**
5.1. Introduction 51

5.2. Service Profile 51
5.2.1. Data Product Specification 52
5.3. Query Model 52
5.4. Service Chain Composition 53
5.5. Register (Meta-model) 53
5.6. Stakeholder/Artefact visibility 54
6. Information Viewpoint - Interoperability Framework 55
6.1. WRON Meta-model 55
6.1.1. Common patterns and components 55
6.1.2. Domain model 56
6.1.3. Metadata model 56
6.1.4. Implementation Model 59
6.2. Modelling methodology 60
6.2.1. Layered models 60
6.2.2. Service Profiles 60
6.2.3. Governance 60
6.2.4. Identity 61
6.2.5. Baseline Information Models 63
6.3. Worked Example 64
6.3.1. Dam level Scenario 64
7. Computational Viewpoint 67
7.1. Overview 67
7.1.1. Extending Best Practice 67
7.1.2. Specification methodology 67
7.1.3. Point-of-Truth vs. Data Warehouse pattern 67
7.2. Key Artefacts (Components) 68
7.2.1. Data Access services 68
7.2.2. Processing Services 69
7.2.3. Simple Access 69
7.2.4. Registry 69
7.2.5. Semantic Registry 69
7.2.6. Supporting Service Types 70
7.3. Computational Framework 70
7.3.1. Role of standards 71
7.3.2. Service Framework 71
7.3.3. Mediation 71
7.3.4. Semantic framework 72
7.3.5. Dissemination and persistence of metadata 72
8. Engineering Viewpoint 74
9. Technology Viewpoint 75
A. ISO 19100 Framework 76
A.1. Overview 76
A.2. Notation - ISO 19103 78
A.3. Feature Model, Feature Type Catalogue-GFM ISO 19109, ISO 19110 79
A.4. Objects and Coverages- ISO 19109, 19123 79
A.5. Utility Types- ISO 19103, 19107, 19108, 19111 80
A.6. Metadata 80
A.7. Registry Model 80
A.8. Profiles 80
A.9. XML Encoding 80
B. ISO 19100 Framework 81
10. Bibliography 84

Executive Summary

The Water Resources Observation Network - Reference Model (WRON-RM) describes a framework for:

- linking Australia's many water and water related data assets;
- harnessing new data streams from satellites and on-ground sensor networks; and
- delivering water information to those who want it and have the right to access it, as and when they need it, in a form that is useful to them.

This framework will enable the coupling of data and processing services to forecasting and reporting technologies which will, in turn, improve the visibility, currency and usability of water information.

The WRON-RM is a tool for system architects, information modellers and system developers working on systems that relate directly or indirectly to water. It allows for the development of systems that can interoperate through distributed services and confirms that the WRON can be viewed as:

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

The WRON-RM is described using the viewpoints of ISO/IEC 10746 Reference Model for Open Distributed Processing (RM-ODP). This describes the WRON Framework from a number of perspectives allowing for greater understanding of the WRON-RM and hence the WRON.

A number of principles have been used in the development of the WRON-RM, the most important of these is 'adopt, adapt, invent'. That is, where appropriate, components of the Reference Model should be adopted from existing work. If this was not appropriate, components should be adapted from existing work. The final resort should be to invent a new component. The reasons for this principle are many. Key amongst them is to ensure close alignment with existing initiatives and hence allow potential interoperation with tools developed within these other frameworks.

The WRON-RM is built around six key system use cases. These are:

1. End User – the process of accessing information through the WRON
2. Data Provision – the process of contributing data to the WRON
3. Functionality Provision – the process of contributing data processing services to the WRON
4. Enablement and Governance – the processes of controlling various components of the WRON
5. Cross-business Domain Integration – the processes enabling integration with other frameworks
6. System Maintenance – the processes associated with maintaining the WRON.

Governance is a very important subject to the WRON-RM and hence the WRON. Lack of appropriate governance regimes for of all components of the WRON poses a considerable risk. Governance regimes will be required for: the Reference Model, Compliance profiles, vocabularies, registries and registers, and data resources.

In developing and describing the WRON Architecture, the following set of principles have been used:

1. Implementation Neutrality – the WRON-RM will not assume particular technologies for implementation. This allows for evolution to occur.

2. Timeliness, Transparency, Traceability – users must have confidence in the information they access though the WRON. This means access must happen in a timely fashion. All information must be traceable to the point of truth and any processing undertaken must be transparent.
3. Scalability – the WRON must have longevity and wide accessibility.
4. Simplicity – all who interact with the WRON should be able to understand it to the level they need to understand it.
5. 'Adequate Description' – descriptions of component behaviour must be adequate for the intended use.
6. Who pays the cost? – efforts in simplification should be targeted at those areas where the most gain is to be made
7. Semantic Robustness – stakeholders must have the ability to declare their adherence to a common understanding of definitions.
8. No Private Contracts – there must be no component of the WRON that depends upon unpublished private agreements which contradict or compromise the published standard.

From an Information perspective, a number of artefacts are defined by the WRON-RM. Of these service profiles and query models are key to ensuring the success of the WRON. These provide for a simple path for the implementation of services and allow for ease of interaction with these services through access to a machine readable description of the functionality supported by individual services.

The interoperability framework of the WRON-RM describes the WRON meta-model and a methodology for creating information elements. This framework provides an approach for creating and the WRON in terms of extended content and increased usability of data.

From a computational perspective the WRON-RM identifies the components: data access services; processing services; simple access services; registries (including semantic registries) and supporting services.

Version 0.1: This version of the WRON-RM is incomplete. It has been developed as an interim step at which external review of the framework can be undertaken. It should, under no circumstances, be viewed as a complete reference model for an interoperability framework.

1. Introduction

1.1. Background

The beginning of the 21st Century has seen renewed focus, both nationally and internationally, on the state of Australia's water resources. This has led to the recognition of a number of challenges that need to be addressed to ensure water security for the nation into the future. These challenges include:

- Ensuring Australia's major cities continue to have water. Current conservative estimates suggest that these cities face a water supply deficit of 850GL by 2030 due to pressures of population growth alone (Water Services Association of Australia 2005);
- Addressing the increasing social and political pressure to deliver on environmental flows to maintain the health of Australia's rivers;
- The Murray Darling Basin is likely to have 2500 GL/year less water flowing into it in 2030 than it does today (van Dijk 2006). This is five times as much water as the 500 GL/year tranche Australia has just spent $1bn to reclaim in the basin via the Living Murray Initiative 'First Step' decision; and
- Understanding the implications of and addressing problems associated with a drying climate.

The purpose of the National Water Initiative (NWI)[1] is to address many of these issues. This is reflected in the NWI's overall objective to achieve a nationally compatible market, regulatory and planning based system of managing surface and groundwater resources for rural and urban use that optimises economic, social and environmental outcomes. The agreement further recognises that achieving this outcome is dependent on focussing on water management.

"You can't manage what you don't measure"

Hon. John Howard 25/1/2007[2]

The NWI states that meeting Australia's water challenges requires excellence in water management. It also requires excellence in water planning, water regulation and an appropriate investment in water infrastructure. Furthermore, it requires all of these things at local, regional and national scales.

Importantly, this excellence can only be achieved through having two key information products:

1. A synoptic understanding of the current status of water; and
2. A reliable forecast as to the probable future status of water.

Without the ability to understand the status and trend of water resources at multiple scales, many of the reforms in water planning, regulation, management and trading which form the blueprint of the National Water Initiative (NWI) can not be realised.

Unfortunately, Australia currently lacks the ability to produce either of these information products. This is because:

[1] http://www.nwc.gov.au/nwi/index.cfm
[2] Address to the National Press Club, http://www.pm.gov.au/media/Speech/2007/speech2341.cfm

- There are gaps in our water data assets – Much of the data that is collected is not collected with the temporal and spatial resolution, currency or accuracy required (Budd et. al. 2004);
- That water data that does exist is often difficult to locate and access - data is spread across more than 200 agencies Australia-wide and does not conform to agreed standards (Budd et. al. 2004);
- The scientific knowledge and tools required for planning and forecasting are underdeveloped and fragmented; and
- No comprehensive reporting tools or systems are in place that allow automation of reporting processes – each time the nation reports on a national, regional or local basis required information flows must be re-invented and the data must be manually collated. The recent AWR2005 Baseline Assessment is an excellent case in point.

1.1.1. The WRON

The Water Resources Observation Network - Reference Model (WRON-RM) describes a framework for:

- linking Australia's many water and water related data assets;
- harnessing new data streams from satellites and on-ground sensor networks; and
- delivering water information to those who want it and have the right to access it, as and when they need it, in a form that is useful to them.

This framework will enable the coupling of data and processing services to forecasting and reporting technologies which will improve the visibility, currency and usability of water resources information. This Water Resources Observation Network (WRON) will allow for improved environmental, social and economic outcomes with respect to water and help to meet the challenges described previously.

Once in place, the WRON will provide the ability to monitor, forecast and manage water demand, supply, quality and usage patterns in any part of the country, at any time. It will provide the transparency and rigour demanded by government, community and business and fundamental to sustainable water resource management.

The concept of 'WRON Compliance' 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. This concept is core to the WRON and confirms that the WRON will be

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

WRON Compliant products and systems will implement some or all of these interfaces and make possible the linking of distributed tools and data sets.

It should be noted that the WRON is not the first project to propose the development of an interoperable, distributed architecture in Australia. It is however, considerably more ambitious than earlier activities such as SeeGRID[3] and SIDP[4], which were discrete technology demonstrators, or the MII[5], which is in an early phase focused on simple catalogues of data

[3] Solid Earth Environment GRID – www.seegrid.csiro.au
[4] Spatial Interoperability Demonstrator Project (SIDP)
[5] Marine Information Infrastructure (aka Oceans Portal)

access services. In particular, the WRON will include Sensor Webs as a form of data input and a greater variety of sophisticated simulation, analysis and reporting tools as clients of the framework. More importantly, it is intended that the WRON will constitute a complete end-to-end information capture, distribution, processing, modelling, reporting and alert system.

1.2. The WRON Reference Model

The WRON Reference Model (WRON-RM) will be a key resource for individuals or groups wishing to develop 'WRON Compliant' tools and, through these tools, contribute to the WRON. The broad subject areas that the WRON-RM addresses are:

1. Information – description of a common, harmonised, extensible model for integration of water information;
2. Distribution and access – identification and definition of the information flows, service interfaces and system components required to support WRON outcomes;
3. Governance – understanding which groups, organisations and individuals are/will be responsible for the various WRON components and how this might affect the WRON architecture;
4. Testing – definition of testing processes and procedures for providing confidence that tools are truly WRON compliant; and
5. Evolution – procedures and infrastructure to support the evolution of WRON to meet emerging challenges, such as external standards, design optimisations and extended information requirements.

Thus, the intended audience for the WRON-RM is system architects, information modellers and system developers working on systems that relate directly or indirectly to water.

1.2.1. Principles of the WRON-RM

The WRON-RM will be developed using the following set of principles:

1. Adopt, Adapt, Invent – 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.
2. Describe only what can be shown to be implementable - the WRON-RM will only describe components that can be implemented and for which a reference implementation exists. The WRON-RM will not describe aspirational components.
3. The WRON-RM will not mandate particular software implementations as part of the solution.

1.2.2. Governance of the WRON-RM

Appropriate governance of key components of the WRON will be vital to ensure its ongoing viability. Similarly governance of the WRON-RM is an issue that must be addressed to ensure success of the WRON.

The WRON-RM must be dynamic. That is, as new concepts and technologies become available, it is essential that the WRON-RM be able to address these and where appropriate,

integrate them. This means that the WRON-RM must be undergoing a process of constant review similar to that undertaken for many international standards.

Ensuring these reviews are undertaken suggests that the WRON-RM must be governed.

Version 0.1: At this point, the governance model for the WRON-RM is yet to be determined. This will be addressed in Version 1.0.

1.2.3. Structure of the WRON-RM

The WRON-RM is described using the viewpoints of ISO/IEC 10746 Reference Model for Open Distributed Processing (RM-ODP). The RM-ODP provides a system description based on five viewpoints, each from a different system perspective. The WRON-RM viewpoints are:

WRON-RM Enterprise Viewpoint – articulates the key business drivers, desired outcomes and adoption strategy for the WRON. It details the relationship between the WRON and other relevant national and international initiatives such as the Australian Water Data Infrastructure Project (AWDIP), National Data Network (NDN) and the European initiative, Orchestra. It provides a description of the environment within which the WRON will be developed, and therefore constraints and obligations (e.g. policies) that must apply in all other viewpoints. The Enterprise Viewpoint specifies the major architectural requirements, which are standards conformant and loosely coupled and developed by the process defined in the WRON Roadmap 1.0. This viewpoint can be found in Section 4.

WRON-RM Information Viewpoint – focuses on the information content of the WRON. The key activities in developing this viewpoint are the identification of information elements, manipulations that may be performed on these elements and information flows. This viewpoint can be found in Sections 5 and 6.

WRON-RM Computational Viewpoint – deals with the logical partitioning of the WRON into functional components independent of any specific environment. A key element of this view is a 'notional architecture'. This is an idealised description of the services and interfaces required to construct the WRON. This viewpoint can be found in Section 7.

WRON-RM Engineering Viewpoint – deals with the practicalities of delivering WRON outcomes from deployed components within a network infrastructure. This viewpoint must provide a framework for assessing the target performance of the system, including predictability of responses, system robustness and processing requirements. This viewpoint can be found in Section 8.

WRON-RM Technology Viewpoint – identifies possible technical artefacts for engineering mechanisms, computational structures, information structures and enterprise structures whilst being as independent of the other four viewpoints as possible. This independence will help to 'future-proof' the WRON. This viewpoint can be found in Section 9.

Figure 1-1: Relationships between RM-ODP Viewpoints

shows the relationship between these viewpoints. It is important to note that the Information and Computational viewpoints inform each other. This is because an understanding of what is 'implementable' is necessary to inform the information model. At the same time, knowledge of the information model is required before developing a service signature. These cross dependencies imply that the two viewpoints must be developed in parallel.

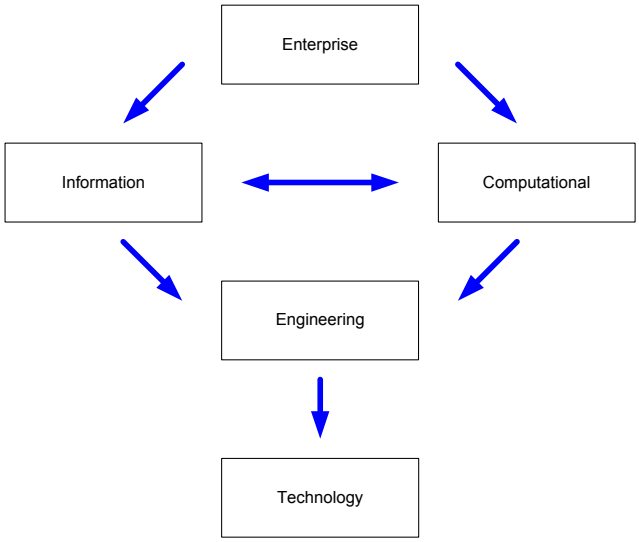

Figure 1-1: Relationships between RM-ODP Viewpoints

1.2.4. Evolution of the WRON-RM

The development of the WRON-RM will be based on a standard, iterative approach to evolve and extend a core model ensuring that implementation aspects are validated; i.e. the reference model can be implemented and is fit for purpose.

It is intended that part of the reference model development process will be an intensive effort of validation and extension which will produce WRON Demonstration systems. The process will be:

1. The System Requirements of the WRON Demonstration System will be captured. This will reflect the use cases of the WRON stakeholders and lead to the identification of Scenarios or Use Cases that define requirements for reference model content.
2. These requirements are translated into proposed or Trial Reference Model Content. Such content should be either a separate identifiable set of modules that extend the WRON-RM, or change requests to modify existing modules to meet an identifiable business need. The trial reference model content may be based on the system design developed for the demonstration system, or the demonstration system may adopt the approach proposed as trial content. To satisfy the need for stakeholder acceptance, it is essential that all trial content be validated via a test implementation before it can be formally mandated in the WRON-RM. This provides a mechanism to deliver research outcomes (WRON Demonstration Systems) into a blueprint for an operational context.
3. Following trial use of content, there is a Post Use Review of Trial Content. This review is based on the experience and knowledge gained from the demonstration system. This review leads either to the trial content progressing to one or more of:
    - adoption of new modules in the WRON-RM;
    - clarification of the initial scenarios/use cases;
    - feedback that leads to changes in the Roadmap – such as adding items to the roadmap, and changing the order in which items are developed; or
    - change request to existing or proposed WRON-RM components, including justification and impact analysis.

Each stage in the reference model development process leads to the publication of a new version of the WRON-RM.

Further detail on the evolution of the WRON-RM is provided in the WRON-RM Roadmap.

1.2.5. Scope of WRON-RM 0.1

Version 0.1 of the WRON-RM is not a complete Reference Model and should not be seen as such. Rather, it has been released on a limited basis for review by peers and potential stakeholders to ensure that the conceptual architecture for the WRON is potentially useful and is, importantly, implementable.

Throughout this document, the reader will find sections for which there is currently no content. These sections will simply state that content will exist for Version 1.0 of the Reference Model.

Delivery of Version 1.0 of the WRON-RM is due towards the end of 2007.

1.3. Context of the WRON-RM

The WRON will be developed and deployed against a background of a number of related information, science and engineering activities both nationally and internationally. These include some specific to water, as well as others concerned more generally with data infrastructures and frameworks.

The following section describes some of the types of initiatives that may influence the development of the WRON-RM. Specific examples of these initiatives are discussed in the WRON-RM Enterprise Viewpoint (Section 4.2.3)

1.3.1. Spatial Data Infrastructures

Spatial data infrastructures (SDIs) concern the deployment of standardised access protocols for geospatially encoded data and services. One of their primary goals is for data sets to be provided ‘live’ from a single point-of-truth, traceable to the (usually) statutory agency with primary responsibility for its maintenance. The desire is to ensure that the best data is always used for analysis and decision support, and to minimise the risks associated with maintaining local copies of critical data sets.

SDI’s often have an emphasis on framework data that is used as input to processes in multiple other agencies. As a means to this end, there is a particular focus on discovery metadata, data portals and standardisation of web interfaces.

The WRON can be viewed as a Spatial Data Infrastructure. Furthermore, it can be seen as an SDI within a network of SDIs. There is a natural hierarchy of systems in practice, with data management devolved to the contributing nodes, coordinated by data standards developed at the aggregating tiers, such as the WRON-RM. It is expected that provision of tools to enable implementation of standards will need to be propagated along with data standards to achieve devolved compliance with the needs of the broader scoped systems.

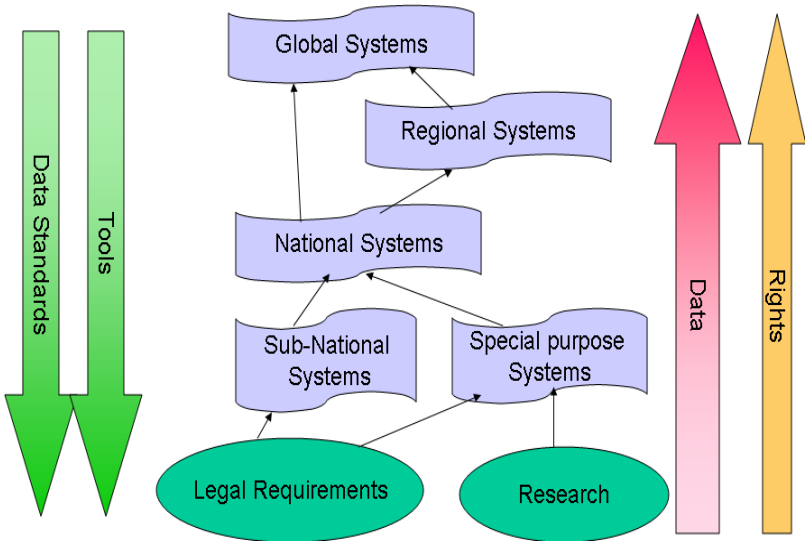

Figure 1-2: Responsibilities in multi-tier systems

1.3.2. GRID Computing

GRID Computing refers to the provision of significant computational and data management resources requiring state as remote commodity services. The development of a standardised suite of tools for GRIDs has moved computing infrastructures from being a narrow interest of ICT researchers and a few communities with specialised requirements, to the attention of the research community in general. 'e-Science' initiatives have received a huge investment by research councils in Europe, and 'Cyberinfrastructure' by NSF in the US. e-Science is also a significant component of several of the Australian NCRIS[6] initiatives.

The standard GRID toolkit has been evolving from Globus toolkit 2 specific services towards convergence with mainstream web-standards, contributing particularly to the WS-XX group of standards, primarily concerning access to stateful resources and authentication management using virtual organisations (Foster, 2006).

1.3.3. Simulation Frameworks

As the WRON-RM is developed, it will be important to be aware of, and consider the variety of, existing simulation frameworks. Of particular interest is the EU's OpenMI[7] which specifies an interface to simulation components that address issues that arise when chaining or cascading models together. Issues such as differing domains, differing spatial and temporal scales and resolutions have been addressed and could be potentially leveraged for the WRON.

Of significant interest and warranting special consideration also, is the modelling framework that is being developed to underpin the eWater CRC[8] products. These products and the associated simulation frameworks will have specific needs that will influence the priority of some of the reference model development.

[6] http://www.ncris.dest.gov.au/
[7] http://www.openmi.org/openminew/
[8] http://www.ewatercrc.com.au/

1.3.4. Water Data

In the water domain, a number of significant projects are underway both nationally and internationally that are of importance to the WRON.

The Australian Water Data Infrastructure Project (AWDIP)[9] is developing a water-based SDI through which various jurisdictional agencies plan to share a defined set of water information.

The Australian Water Resources Information System (AWRIS) is being developed by the National Water Commission (NWC) to provide a suite of tools for analysing and visualising water related data. It is intended that these tools will use an extended version of the SDI being developed by AWDIP to access this data.

In the United States, the Consortium of Universities for the Advancement of Hydrologic Science (CUASHI)[10] is developing a Hydrologic Information System (HIS) for distributing water related information.

These projects are described in more detail in Section 4.2.3.

1.4. The Dam Level Scenario

Throughout this document examples are given to help describe the concepts being presented. These examples are clearly differentiated from the body of the document and all relate to a single scenario – the Dam Level Scenario.

The Dam Level Scenario consists of a set of key roles and use cases associated with the provision and use of historical, current and future dam level information. The key roles for the Dam Level Scenario are described in Table 1-1 along with an indication of how these specific roles map to WRON System Roles.

Table 1-1: Dam Level Scenario Roles and associated WRON System Roles

Role	Description	Relationship to WRON System Roles
End User	An individual or group interested in dam level information at the single dam or water region level.	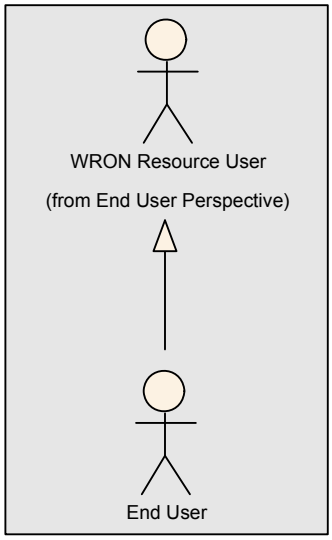

[9] http://www.daffa.gov.au/brs/water-sciences/ground-surface/awdi-project
[10] http://www.cuahsi.org/

Water Index Service Provider	An individual or group that provides a service that generates a water index given a dam level.

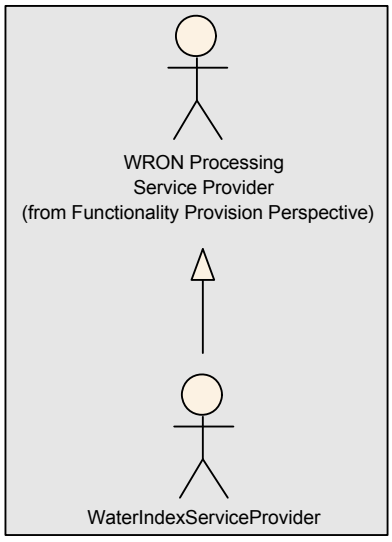

Water Index Definer	An individual or group who defines the concept of a water index
Water Region Definer	An individual or group responsible for defining water regions. This might be a State Water Manager.

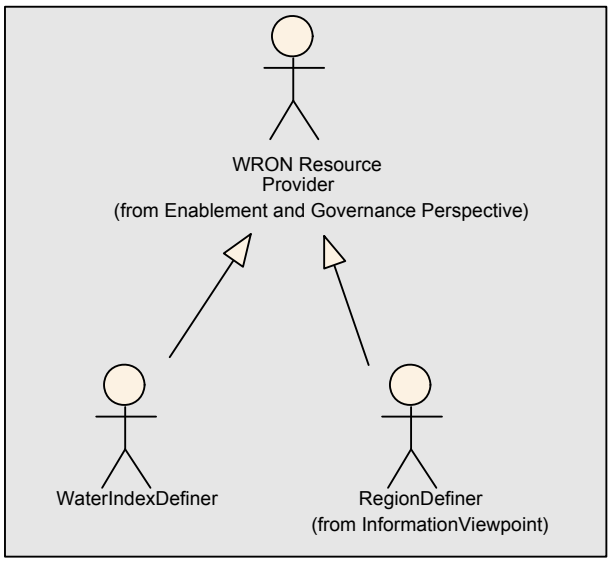

Dam Level Provider	An individual or group responsible for providing dam level information for a specific dam

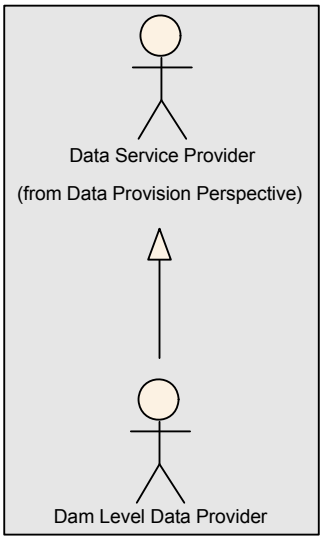

Although very simple, the Dam Level Scenario contains representatives of all system use cases that need to be supported by the WRON (and hence the WRON-RM). It is hoped that, by using a common scenario throughout the WRON-RM, readers will be able to see how the various RM--ODP viewpoints inter-relate.

2. Notation

2.1. Glossary

2.1.1. Abbreviations

Abbreviation	Meaning
AAAA	Authentication, Authorisation, Accounting and Audit
API	Application Program Interface
AWDIP	Australian Water Data Infrastructure Project
AWRIS	Australian Water Resource Information System
CRS	Coordinate Reference System [ISO 19111]
CSIRO	Commonwealth Scientific and Industrial Research Organisation
CSL	Conceptual Schema Language [ISO 19103]
DDL	Data Definition Language [SQL]
EB-XML	Electronic Business Using eXtensible Markup Language
ebRIM	EB-XML Registry Information Model
ESCAWRI	Executive Steering Committee for Australia's Water Resource Information
FTC	Feature Type Catalogue [ISO 19110]
GFM	General Feature Model [ISO 19109]
GI	Geographic Information
GIS	Geographic Information System
GML	Geography Markup Language [OGC, ISO 19136]
ID	Identifier
INSPIRE	Infrastructure for Spatial Information in Europe
ISO	International Organization for Standardization
ISO/TC 211	ISO Technical Committee 211 – Geographic Information
NWI	National Water Initiative
OASIS	Organisation for the Advancement of Structured Information Standards http://www.oasis-open.org/who/
OGC	Open Geospatial Consortium (formerly Open GIS Consortium)
OMG	Object Management Group
ORCHESTRA	Open Architecture and Spatial Data Infrastructure for Risk Management (European Commission IST project)
OWL	Web Ontology Language
OWL-S	Web service ontology based on OWL
RM-OA	Reference Model for the Orchestra Architecture

RM-ODP	Reference Model for Distributed Processing
EV	Enterprise Viewpoint (RM-ODP)
IV	Information Viewpoint (RM-ODP)
CV	Computational (Service) Viewpoint (RM-ODP)
Eng	Engineering Viewpoint (RM-ODP)
TV	Technology Viewpoint (RM-ODP)
SDI	Spatial Data Infrastructure
SOA	Service Oriented Architecture. A service-oriented architecture is a collection of services that communicate with each other. The services are self-contained and do not depend on the context or state of the other service. They work within a distributed systems architecture.
SOAP	Simple Object Access Protocol (W3c) http://www.w3.org/tr/2000/note-soap-20000508/
SWE	Sensor Web Enablement [OGC]
UDDI	Universal Description, Discovery and Integration
UML	Unified Modeling Language [OMG]
URI	Uniform Resource Identifier
URL	Uniform Resource Locator
W3C	World Wide Web Consortium
WRON	Water Resources Observation Network
WSDL	Web Services Description Language (W3C) http://www.w3.org/tr/wsdl
XML	Extensible Mark-Up Language (W3C) http://www.w3.org/xml/
XSD	XML Schema Definition

2.1.2. Terms and Definitions

Term	Definition
Accounting [RM-OA]	Process of gathering information about the usage of resources by subjects.
Application [RM-OA]	Use of capabilities, including hardware, software and data, provided by an information specific system to the satisfaction of a set of user requirements in a given application domain.
Application Domain [RM-OA]	Integrated set of problems, terms, information and tasks of a specific thematic domain that an application has to cope with.
Application Schema [RM-OA]	Conceptual schema for data required by one or more applications.
Architecture	The organisational structure and operating environment of the SDI, including the relationships between its parts, and the principles and guidelines governing their design and evolution

	over time.
Authentication [RM-OA]	Process of verifying the principal of a certain subject. In other words authentication indicates whether a subject is allowed to use a certain principal.
Authorisation [RM-OA]	Process of determining whether a subject is allowed to have the specified types of access to a resource. Usually, authorisation is carried out in the context of authentication. Once a subject is authenticated it may be authorised to perform different types of access.
Binding	Specific syntax and parameter values used by a client to invoke a specific server operation.
Catalogue [RM-OA]	A registry that, in the SDI context, is usually used to describe spatial data sets. Collection of entries, each of which describes and points to a feature collection. A catalogue registers the existence, location and description of feature collections held by an Information Community.
Client	A software component or an application that accesses a service. Clients may be categorised in three ways Thin clients where the client supports only human-interface code, such as a web browser or a minimal pda or wifi handset, and must also support non-proprietary standards. They typically lack long-term memory such as disk drives. Application code and data access both run remotely and are entirely dependent on an external network connection. Thick clients where the client supports all the human interface and application code, may support some or all data access code, and may support long-term data memory. Human interface code may be entirely customised and not conform to non-proprietary standards. May not even support human interfaces i.e. May be entirely automated remote processes. May operate at times without network connection. Chubby clients have capabilities somewhere on the spectrum between thick and thin clients i.e. May support some application and data code, and may store limited amounts of data. Will usually but not necessarily support human interfaces. May operate well for limited time without network connection.
Conceptual architecture	An overview of the services, data, technology and institutional environment of SDI. It describes, in general terms, both what the SDI will include and how it will operate.
Conceptual Model [RM-OA]	Model that defines concepts of a universe of discourse.
Conceptual Schema [RM-OA]	Formal description of a conceptual model.
Contract [based on http://en.wikipedia.org/wiki/Design_by_contract]	Formalisation of obligations and benefits between elements of a software system or between stakeholders and the system.
Coverage [ISO	Function from a spatial, temporal or spatiotemporal domain to

19123]	an attribute range. A coverage associates a position within its domain to a record of values of defined data types. Thus, a coverage is a feature with multiple values for each attribute type, where each direct position within the geometric representation of the feature has a single value for each attribute type.
Custodian	The authoritative manager of an SDI resource, whether data set, service or component, who is responsible for the declaration of the policies regarding use and accounting for the resource.
Datastore	Any type of persistent storage for components and data. Content may be static or dynamic. May include database systems, file systems, structured text storage, xml repositories etc.
Discovery [RM-OA]	Act of locating a machine-processable description of a resource that may have been previously unknown and that meets certain functional criteria. It involves matching a set of functional and other criteria with a set of resource descriptions.
Engineering Viewpoint	Viewpoint of the WRON Reference Model that specifies the practicalities of delivering WRON outcomes from deployed components within a network infrastructure.
End User	Members of organisations, government or business that are involved in an application domain (e.g. water resources) and that use the applications built as part of the WRON.
Enterprise Viewpoint	Viewpoint of the WRON Reference Model that specifies the business drivers, outcomes and adoption strategy for the WRON.
Feature [derived from ISO 19101]	Abstraction of a real world phenomenon [ISO 19101] perceived in the context of a WRON application.
Feature Catalogue [http://www.opengeospatial.org/ resource/glossary]	Catalogue containing definitions and descriptions of the feature types, feature attributes, and feature relationships occurring in one or more sets of geographic data, together with any feature operations that may be applied
Feature Collection [http://www.opengeospatial.org/ resource/glossary]	A special category of feature that represents a collection of features that have common metadata and formal relationships. "A set of related features managed as a group. Feature collections can be identified at different abstraction levels, i.e. high abstraction level, e.g. ‘topography’ and low abstraction level, e.g. ‘roads’. The terms feature, feature collection and coverage are defined in line with OpenGIS 5."
Framework [http://www.opengeospatial.org/ resource/glossary]	An information architecture. In terms of software design, a reusable software template, or skeleton, from which key enabling and supporting services can be selected, configured and integrated with application code.
General Feature Model [http://www.opengeo	Metamodel of feature types. A feature may have properties that may be operations, attributes or associations. Any feature may have a number of attributes, some of which may be geometric and spatial. A feature is not defined in terms of

spatial.org/ resource/glossary]	a single geometry, but rather as a conceptually meaningful object within a particular domain of discourse, one or more of whose properties may be geometric.
Geospatial [http://www.opengeospatial.org/ resource/glossary]	Referring to location relative to the Earth's surface. 'Geospatial' is more precise in many GI contexts than 'geographic', because geospatial information is often used in ways that do not involve a graphic representation, or map, of the information.
Implementation [http://www.opengeospatial.org/ resource/glossary]	A software package that conforms to a standard or specification. A specific instance of a more generally defined system.
Information Community [http://www.opengeospatial.org/ resource/glossary]	A collection of people (a government agency or group of agencies, a profession, a group of researchers in the same discipline, corporate partners cooperating on a project, etc.) who, at least part of the time, share a common digital geographic information language and common spatial feature definitions.
Information Viewpoint	Viewpoint of the WRON Reference Model that specifies the information model for the WRON.
Interface [RM-OA]	Named set of operations that characterize the behaviour of an entity. The aggregation of operations in an interface, and the definition of interface, shall be for the purpose of software reusability. The specification of an interface shall include a static portion that includes definition of the operations. The specification of an interface shall include a dynamic portion that includes any restrictions on the order of invoking the operations.
Interoperability [RM-OA]	Capability to communicate, execute programs, or transfer data among various functional units in a manner that requires the user to have little or no knowledge of the unique characteristics of those units [ISO 2382-1].
Meta-information [RM-OA]	Descriptive information about resources in the universe of discourse. Its structure is given by a meta-information model depending on a particular purpose. Note: A resource by itself does not necessarily need meta-information. The need for meta-information arises from additional tasks or a particular purpose (like catalogue organisation), where many different resources (services and data objects) must be handled by common methods and therefore have to have/get common attributes and descriptions (like a location or the classification of a book in a library).
Meta-information Model [RM-OA]	Implementation of a conceptual model for meta-information.
Middleware [http://www.opengeospatial.org/ resource/glossary]	Software in a distributed computing environment that mediates between clients and servers.

Ontology [RM-OA]	Explicit, formal specification of a shared conceptualisation (Gruber 1993). It is formal in order to not only make it readable by humans, but also by machines. It is explicit as it is based on a taxonomy specified in terms of concepts, properties (or relationships) and axioms (the vocabulary). It is shared in the sense that these specifications are fixed as an agreement set up and shared by a dedicated user community and that it is associated with a particular subject area (domain) or task. It is a conceptualisation as it defines a conceptual schema by abstracting from a real or hypothetical world. Its ultimate purpose is to enable machine understanding which in turn provides potential for data and service interoperability.
Open Architecture [RM-OA]	Architecture whose specifications are published and made freely available to interested vendors and users with a view of widespread adoption of the architecture. An open architecture makes use of existing standards where appropriate and possible and otherwise contributes to the evolution of relevant new standards.
Operation [RM-OA]	Specification of a transformation or query that an object may be called to execute. An operation has a name and a list of parameters.
Purpose (of meta-information) [RM-OA]	A purpose of meta-information describes the goal of the usage of the resources.
Reference Model [RM-OA]	A reference model is a framework for understanding significant relationships among the entities of some environment, and for the development of consistent standards or specifications supporting that environment. A reference model is based on a small number of unifying concepts and may be used as a basis for education and explaining standards to a non-specialist.
Registry	A listing of the specific, individual services, components, datasets or other entities that comprise the SDI or are relevant to its users. Instance registries are used to identify, locate, and describe individual instances. Many registries refer to associated type libraries that record the allowed types within registry classes e.g. Types of services, types of user authorities.
Resource	Data, services and components that are published and underlie the creation of all useful products. Resources are presented to the internet as web services.
Schema	A schema is an expression of the type using a particular data modelling language. Types can be described as classification taxonomy for a set of schema definitions. The OGC application data modelling language is GML and each schema fragment corresponding to a given type is defined in GML.
Semantic Interoperability [RM-OA]	Semantic interoperability emphasizes the importance of information inside enterprise networks and focuses on enabling content, data, and information to interoperate with software systems outside their origin. Information's meaning is the crucial enabler that allows software to interpret the

	appropriate context, structure, and format in which the information should reside at any given moment and inside any given system.
Semantic Web [RM-OA]	The Semantic Web provides a common framework that allows data to be shared and reused across application, enterprise, and community boundaries. It is a collaborative effort led by W3C with participation from a large number of researchers and industrial partners. It is based on the Resource Description Framework (RDF), which integrates a variety of applications using XML for syntax and URIs for naming.
Service	A collection of operations, accessible through one or more interfaces, that allows a user to evoke behaviour of value to that user. A server delivers each service. A service may encapsulate many processes. A 'service instance' is another name for a server (b).
Service Profile	
Spatial Data Infrastructure [RM-OA]	Relevant base collection of technologies, policies and institutional arrangements that facilitate the availability of and access to spatial data. The Spatial Data Infrastructure 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.
System [RM-OA]	Something of interest as a whole or as comprised of parts. Therefore a system may be referred to as an entity. A component of a system may itself be a system, in which case it may be called a subsystem. Note: For modelling purposes, the concept of system is understood in its general, system-theoretic sense. The term 'system' can refer to an information processing system but can also be applied more generally.
System User [RM-OA]	Provider of services that are used for an application domain as well as IT architects, system developers, integrators and administrators that conceive, develop, deploy and run applications for an application domain.
Technology Viewpoint	Viewpoint of the WRON Reference Model that specifies the technological choices of the service platform and its operational issues.
Transaction [RM-OA]	Transaction is a feature of the architecture that supports the coordination of results or operations on state in a multi-step interaction. The fundamental characteristic of a transaction is the ability to join multiple actions into the same unit of work, such that the actions either succeed or fail as a unit.
Universe of Discourse [ISO 19101]	View of the real or hypothetical world that includes everything of interest.
User	Human acting in the role of a system user or end user of the WRON architecture.
Viewpoint [RM-ODP]	Subdivision of the specification of a complete system, established to bring together those pieces of information

	relevant to some particular area of concern during the design of the system.
Web Coverage Service (WCS)	Supports electronic interchange of geospatial data as ‘coverages’ – that is, digital geospatial information representing space-varying phenomena. A WCS provides access to potentially detailed and rich sets of geospatial information, in forms that are useful for client-side rendering, multi-valued coverages and input into scientific models and other clients. The WCS may be compared to the OGC web map service (WMS) and the ogc web feature service (WFS); like them it allows clients to choose portions of a server's information holdings based on spatial constraints and other criteria. Unlike WMS (OGC Document 01-068R3), which filters and portrays spatial data to return static maps (rendered as pictures by the server), the web coverage service provides available data together with their detailed descriptions; allows complex queries against these data; and returns data with its original semantics (instead of pictures) which can be interpreted, extrapolated, etc. - and not just portrayed. Unlike WFS (OGC Document 02-058), which returns discrete geospatial features, the web coverage service returns representations of space-varying phenomena that relate a spatio-temporal domain to a (possibly multidimensional) range of properties.
Web Feature Service (WFS)	Serves vector data (points, lines and polygons) to the web for use by applications on remote websites. Provides interfaces for describing data manipulation operations on geographic features using http as the distributed computing platform. A web feature service request consists of a description of query or data transformation operations that are to be applied to one or more features. The request is generated on the client and is posted to a web feature server via http. The web feature server then reads and (in a sense) executes the request. The OGC web map service (WMS) allows a client to overlay map images for display served from multiple web map services on the internet. In a similar fashion, the OGC web feature service allows a client to retrieve geospatial data encoded in geography markup language (GML) from multiple web feature services
Web Map Service (WMS)	Produces maps of georeferenced data. A map is a visual representation of geodata; a map is not the data itself. These map views are rendered in a 2d pictorial format such as png, gif or jpeg. The WMS specification thus enables the creation of a network of distributed map servers from which clients can build customised maps. A particular WMS provider in a distributed WMS network need only be the steward of its own data collection. This stands in contrast to vertically integrated web mapping sites that gather in one place all of the data to be made accessible by their own private interface.
Web Service	Application logic accessible across a network using standard internet protocols. Web services combine the best aspects of component-based development and the web. Like components, web services represent functionality that can be easily reused without knowing how the service is implemented. Unlike current component technologies that are

accessed via proprietary protocols, web services are accessed via ubiquitous web protocols (e.g. HTTP) using universally accepted data formats (e.g. XML).

2.2. UML Diagrams

Data model and use case diagrams that appear in this document are presented in accordance with the Unified Modelling Language (UML) specified in ISO/IEC 19501-1. The UML notation for these diagrams is described below.

2.2.1. Use Cases

The Use Case diagrams included in this document use the UML notation as shown in Figure 2-1.

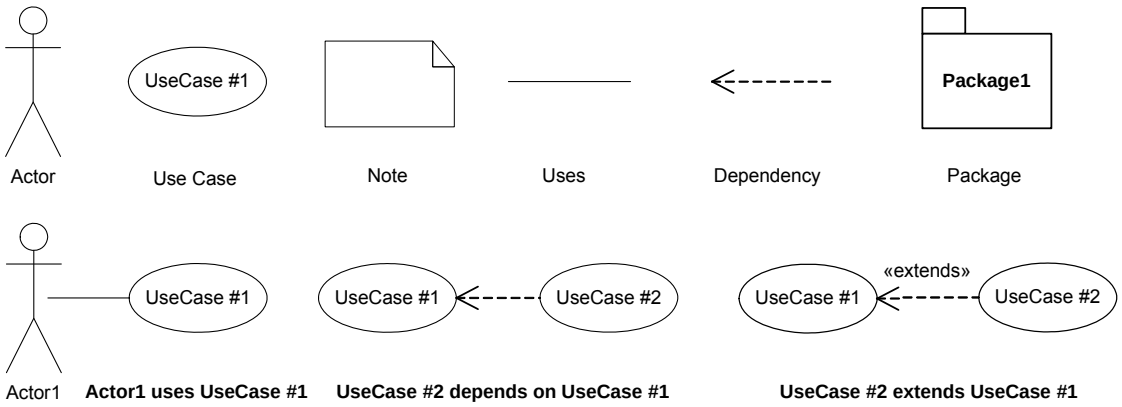

Figure 2-1: UML use case diagram notation

2.2.2. Data Models (Class Diagrams)

Figure 2-2 shows the UML notation used within the data model diagrams in this document. UML Class Diagram notation is used.

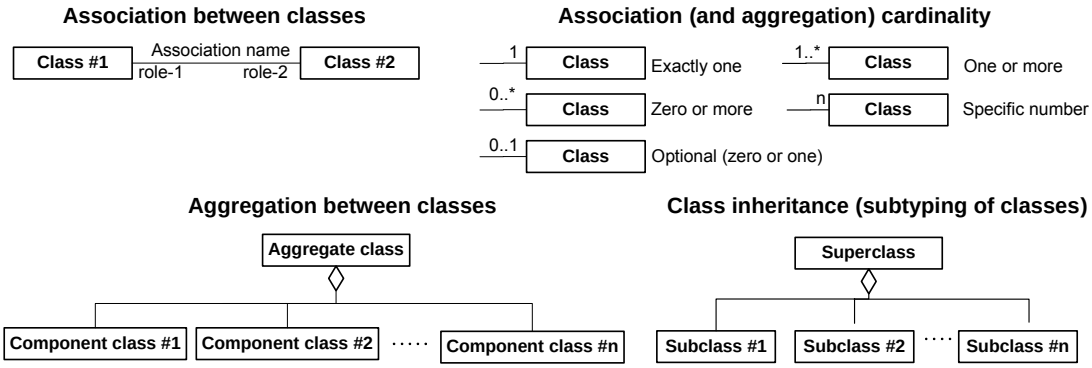

Figure 2-2: UML notation for class diagrams[11]

[11] Extract from OGC Web Services Common Specification (Figure 1)

3. References

The following references are used as background documents for the WRON-RM.

3.1. ISO Standards (or Drafts)

ISO/IEC 10746-1:1998 (E). Information technology – Open Distributed Processing – Reference model.

ISO 19101:2002 (E). Geographic information – Reference model

ISO/AS/NZS 19103:2006. Geographic information – Conceptual schema language

ISO 19107:2004 (E). Geographic information – Spatial schema

ISO 19109:2005. Geographic information – Rules for application schema

ISO 19110:2005. Geographic information – Methodology for feature cataloguing

ISO 19115:2003. Geographic information – Metadata

ISO 19135:2005. Geographic information – Procedures for item registration

3.2. OGC Best Practice Documents

OGC 05-087r4:2006. Observations and Measurements

3.3. Documents and Books

ORCH-D3.2.3 (2007). Reference Model for the ORCHESTRA Architecture (RM-OA). Deliverable D3.2.3. Integrated Project 511678 ORCHESTRA. Publisher: ORCHESTRA Consortium. Version 2.0. 31 January 2007

4. Enterprise Viewpoint

4.1. Overview

The Enterprise Viewpoint of the WRON-RM describes the business purpose for the WRON, from a system behaviour perspective, and the components that must be implemented for it to provide this behaviour. Section 4.2 establishes the context of the WRON-RM in relation to the WRON itself, and to other existing initiatives and systems. In Section 4.3, the key architectural principles and requirements underlying development of the WRON are identified and discussed. A clear governance strategy is vital to the success and adoption of the WRON, and is discussed in Section 4.4.

The final section of the Enterprise Viewpoint (Section 4.5) identifies the principal system use cases for the WRON architecture. This section is organised into perspectives corresponding to particular stakeholders and their need to understand certain aspects of the operation of the WRON. The intent of this approach is to ensure that individual stakeholders can readily identify the aspects of the WRON that they will need to interact with.

In the viewpoints following the Enterprise Viewpoint, the roles, components and information artefacts required to meet the requirements described here will be further defined. Specifically, the Information Viewpoint will identify the nature of the information required to realise the WRON as an information distribution framework. The Computational Viewpoint will define the components that users will expect to interact with and how these relate to each other. In the Engineering Viewpoint, the business needs identified here will be transformed into an engineering solution design that addresses issues such as deployment, size and robustness. Finally, the Technical Viewpoint will identify the technical issues that must be validated in order to implement a WRON architecture that is flexible enough to allow multiple technical platforms to co-exist and evolve. This is necessary as it is recognised that the WRON must provide a viable solution for the existing technical platforms and skills capacity of the various stakeholder communities.

4.2. Business Requirements for the WRON

A key element of the Enterprise Viewpoint is to describe the business environment in which the WRON will exist and, in particular, the elements of this environment that have a direct effect on architectural choices.

The key business requirement requiring development of the WRON was introduced in Section 1. The requirement is to support the need for:

1. A synoptic understanding of the current status of water; and
2. A reliable forecast as to the probable future status of water.

From a systems architecture view, this requirement can be addressed in a number of ways. Examples include:

- Creation of a single organisation which has responsibility for all aspects of water management throughout the nation. This organisation would be responsible for the capture of all water information and can implement solutions that do not require giving access to this information outside the organisation
- Development of a centralised water data warehouse into which water information is copied from other data custodians.
- Development of a distributed environment whereby individual organisations are responsible/encouraged to provide access to the information they are responsible for.

The architecture choices prescribed by the WRON-RM have been determined by investigating the WRON Business environment in three ways:

1. Risk Mitigation – which solution addresses the most risks associated with development of the WRON;
2. Contribution to WRON objectives – which solution will be able to meet the various objectives of the entire WRON initiative; and
3. Alignment of other initiatives – which solution aligns most closely with similar initiatives both nationally and internationally.

4.2.1. Risk Mitigation

An important way of approaching the design of the WRON is to look at the key implementation risks and hence the architectural requirements that are needed to meet them. Table 4-1 lists these risks and describes how the proposed architecture meets them.

Table 4-1: Risks to be addressed by WRON-RM

Risk	Response(s)
Failure to Function (WRON does not provide useful information to stakeholders because of lack of functionality)	Flexibility of functionality through composition of services. Flexibility of ability to deliver data to external platforms and frameworks. Scenario driven testing for initial implementation. Analysis of typical cases for addition of functionality. Implementation Neutrality
Lack of data (stakeholders do not expose sufficient data to achieve the WRON's aims)	Governance: Accountability Demonstrable feasibility Provision of reference implementations Regulatory or financial incentives
Poor data integratability	Common data products Data Model coherence Service Profiles with compliance tests
Timeliness of data delivery (intermediate processing)	Avoid ad-hoc data management between custodian and user Automation requirements built into service profiles and data products Distributed point-of-truth to simplify each component
Robust delivery	Enable data warehousing as automated 'forward-caches' to provide high-availability services with impacting governance and data management scalability.

4.2.2. Contribution to the WRON's Objectives

The goal of the WRON is to establish the technology platform to provide an Australia wide network of water information systems delivering dynamic, timely reporting and forecasting of Australia's water resources with the aim of enabling the reforms in water planning, regulation, management and trading that form the blueprint of the National Water Initiative.

The WRON project seeks to achieve this goal by focusing on four key objectives:

- Improving the usability and availability of water data;
- Enabling water information interoperability;
- Development of modelling and forecasting tools; and
- Development of operational reporting and visualisation tools.

To meet these objectives, five areas of research have been identified – Sensorization, Data Integration, Forecasting Systems, Reporting Systems and Standards & Access (see Figure 4-1).

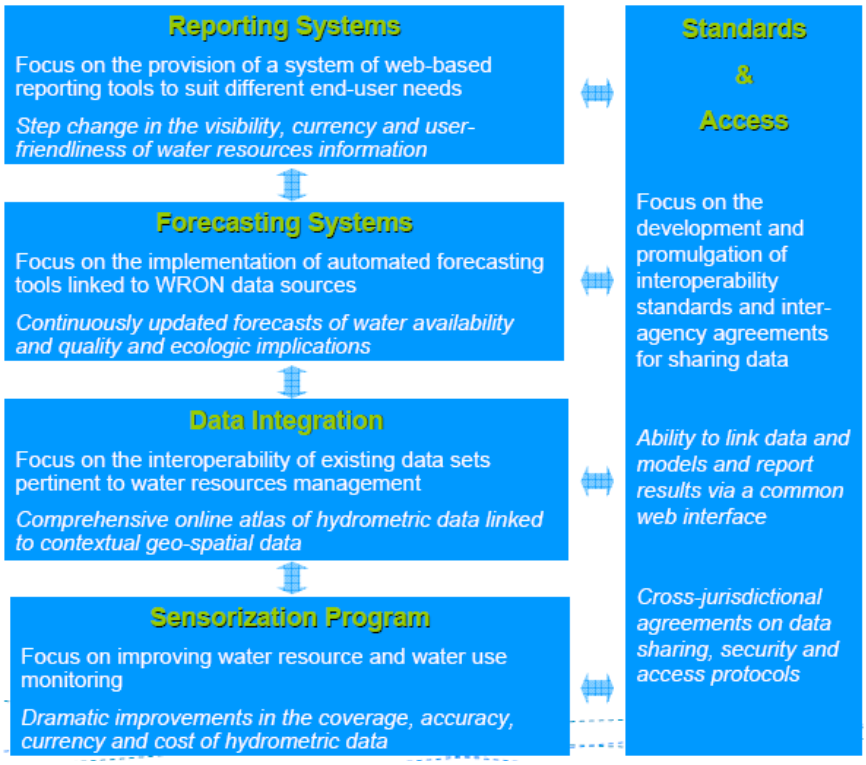

Figure 4-1: Key elements of the WRON

The WRON Reference Model (WRON-RM) contributes to the WRON by fulfilling the first of the WRON's objectives:

> *Enable water information interoperability – through research investments in standards development, web service integration, semantic web, model interoperability and a close partnership with State and Commonwealth entities such as the National Water Commission (NWC) to deliver these standards operationally across Australia (CSIRO, 2006).*

The WRON-RM defines a framework for interoperability enabling the coupling of data and processing services to forecasting and reporting technologies. This framework also provides a platform for the creation of standards and agreements for data sharing and use which will underpin water resource data integration.

The concept of WRON Compliance underlies the WRON-RM and declares the goal of building the WRON from independent components under the management and control of various agencies, built on standard, open interfaces defined by the WRON-RM. The WRON-RM framework provides an architecture for a set of components brought together as a loosely coupled system based on distributed services conforming to standard interfaces.

4.2.3. Alignment with Other Initiatives

Section 1.3 identified the different types of initiatives that may influence the development of the WRON-RM. These are:

1. Spatial Data Infrastructures;
2. GRID Computing;
3. Simulation Frameworks; and
4. Water related projects.

The following section describes specific initiatives that have been identified as being of interest to the WRON-RM

CUAHSI Hydrologic Information System (HIS)

In the United States, the Consortium of Universities for the Advancement of Hydrologic Science (CUASHI) is developing the Hydrologic Information System (HIS) for distributing water related information.

The stated goals of this initiative are to:

> *unite the nation's (USA) water information, to make it universally accessible and useful, and to provide access to the data sources, tools and models that enable the synthesis, visualization and evaluation of the behavior of hydrologic systems.*[12]

At the implementation level, the HIS uses SOAP services to distribute information and hence is relatively simple to implement. Furthermore, underpinning the HIS is the recently released Observations Data Model (ODM)[13] and associated Water Markup Language (WaterML)[14].

From the perspective of the WRON-RM, keen attention should be paid to the further development of WaterML as well as the rate at which the HIS Services are adopted. CUAHSI has a close relationship with the vendor of a widely used software product and hence, it is envisaged a number of the users of this product may adopt some or all of the HIS solution.

Australian Water Data Infrastructure Project (AWDIP)

The Australian Water Data Infrastructure Project (AWDIP) is developing a water domain based SDI through which various jurisdictional agencies plan to share a defined set of water information. The stated business case for AWDIP is to "facilitate national assessments of Australia's water resources through the ongoing development of a comprehensive and accessible national water information framework"[15].

The project is managed by the Executive Steering Committee for Water Resource Information (ESCAWRI) which is constituted under the Council of Australian Governments (COAG) with joint approval by all states. As such AWDIP represents the only national data infrastructure in Australia, with shared responsibilities by all states.

To date AWDIP has defined a specification for data services to provide access to seven water quality parameters and implemented part of this using the OGC Web Feature Service specification using the GeoServer WFS server.

---

[12] http://www.cuahsi.org/his.html
[13] Released March 2007
[14] Released March 2007
[15] ESCAWRI Membership Terms of Reference (available online from: http://www.daffa.gov.au/brs/water-sciences/ground-surface/awdi-project).

AWDIP is seen as an excellent first step in the development of the WRON. The WRON-RM needs to ensure that the achievements of AWDIP are aligned with the WRON.

Australian Water Resources Information System (AWRIS)

In 2005 the National Water Commission (NWC) commissioned the Australian Water Resources 2005 (AWR2005) Baseline Assessment project (NWC, 2005). A key output of the project was a set of specifications for a suite of tools that will deliver future assessments of Australia's water resources. This suite of tools is to be known as the Australian Water Resources Information System (AWRIS).

Fundamental to AWRIS is the recognition that it is not cost effective to undertake assessments (such as AWR2005) using a process of manual data identification, collection and analysis. Rather assessments need to be delivered through toolset such as AWRIS accessing the information required via an underlying enabling framework.

The enabling framework described in the AWRIS specifications is intended to be an enhancement on framework being developed by AWDIP (described above). For this reason, the WRON-RM should carefully consider the requirements of this suite of tools.

National Data Network

The National Data Network (NDN) is designed to provide access to 'a distributed library of data holdings relevant to policy analysis and research'[16]. Development is currently being funded by the Australian Bureau of Statistics (ABS).

Data holdings available through the NDN remain in the control of the Custodian organisations. Discovery of these data holdings is achieved through a central catalogue of all available services. Searching is performed through the central NDN website.

A key element of NDN is the central 'hub', called NDN Central, and its links to the various NDN nodes. These nodes are associated with individual agencies and contain metadata describing all of that agency's data holdings.

Of interest to WRON-RM is NDN's handling of Authentication, Authorisation and Auditing issues. These same issues will need to be addressed for WRON.

ORCHESTRA/INSPIRE

ORCHESTRA is designing and implementing specifications for a service oriented spatial data infrastructure for improved interoperability among risk management authorities in Europe. The ORCHESTRA Architecture is open, based on standards and described in the document Reference Model–ORCHESTRA Architecture (RM-OA).[17]

This project is of considerable interest to WRON-RM as both projects share a number of architectural principles. In particular, their reliance on service oriented architecture and the need to support a distributed environment.

ORCHESTRA is one of a number of groups involved in the development of the INfrastructure for SPatial InfoRmation in Europe (INSPIRE) initiative. This, much broader, initiative aims to create a European spatial information infrastructure that delivers to the users integrated spatial information services. These services will allow users to identify and access spatial or geographical information from a wide range of sources, from the local level to the global level, in an inter-operable way for a variety of uses.[18]

---

[16] http://www.nationaldatanetwork.org/NDN/NDNHome.nsf/Home/About%20Us
[17] http://www.eu-orchestra.org/overview.shtml
[18] http://inspire.jrc.it/home.html

HarmonIT/OpenMI

The Open Modelling Interface and Environment (OpenMI) has been developed by the HarmonIT project in Europe. Its purpose is to allow the seamless linking of individual models to construct 'whole of catchment'' models in order to support integrated catchment management.[19]

OpenMI, like WRON, concentrates on the definition of standard interfaces. In the case of OpenMI, the interface has three functions:[20]

1. *Model definition*: To allow other linkable components to discover items a particular model can exchange in terms of quantities simulated and the locations at which the quantities are simulated.
2. *Configuration*: To define what will be exchanged when two models have been linked for a specific purpose.
3. *Runtime operation*: To enable the model to accept or provide data at run time.

The development of WRON-RM needs to follow closely the progress of this project to determine if there is a need to support this interface.

Catchment Modelling Toolkit

The Catchment Modelling Toolkit (the Toolkit) is a repository of tools intended to improve the efficiency and standard of catchment modelling.[21] The initiative was begun by the CRC for Catchment Hydrology in 1999 and is now sponsored by the eWater CRC.

From the perspective of the WRON-RM, the Toolkit forms one of the key sources of use cases around which the WRON-RM will be developed. eWater CRC sponsors both the Toolkit as well as the initial project to define the WRON-RM. There is also some level of commitment to ensure that eWater products are 'WRON Compliant'. It can be envisaged that there will be a need to develop WRON interfaces for many of the tools in the Toolkit.

NEON/CLEANER

The National Ecological Observatory Network (NEON)[22] and the Collaborative Large-scale Engineering Analysis Network for Environmental Research (CLEANER)[23] initiatives in the US are also of interest to the WRON and hence the development of the WRON-RM.

In particular, both these initiatives propose the development of a large scale sensor networks deployed throughout catchments and the creation of virtual data repositories accessible by many.

4.3. Architectural Requirements and Principles for the WRON

There are a number of design constraints on the WRON system that arise from the broader context within which the WRON must operate over a considerable period of time. These requirements stem from the necessity to support multiple stakeholders and the need for a flexible infrastructure capable of addressing emerging issues within the domain and exploiting new approaches, tools and data as they become available.

---

[19] http://www.harmonit.org/overview/overview.htm
[20] http://www.openmi.org/documents/A_OpenMI_Scope.pdf
[21] http://www.toolkit.net.au/cgi-bin/WebObjects/toolkit
[22] http://www.neoninc.org/about/
[23] http://cleaner.nacse.org/about/index.html

This long-term view characterises Open Distributed Systems – which can be upgraded or extended component by component and consequently has a lifetime expectancy at least an order of magnitude greater than individual technologies. This contrasts with 'stovepipe' or 'warehouse' approaches that tend to live only as long as the shortest lived of the technology platform, resourcing for data management processes, governance arrangements and client needs for centralised access.

The following section describes a set of architectural requirements and principles that the WRON-RM needs to adhere to.

4.3.1. Implementation Neutrality

The WRON-RM will not assume particular technologies for implementation, but rather assume that implementations will change and evolve over time, and that the WRON must continue to function.

This technology neutrality applies to both the particular components that need to be established and also the technical standards, such as protocols (e.g. OGC, W3C) which can be seen to be a technology platform in itself.

That said, the WRON-RM must be proven to be feasible with currently available platforms, without assuming these will be the long term solution.

4.3.2. Timeliness, Transparency, Traceability

A key requirement is for the WRON to provide a level of confidence to the End User and the Data Provider in the integrity of data provision process.

Accordingly, it is important for data to be traceable to the 'point of truth', or authoritative data provision, even if this is not the place the user gets the final packaged data product.

This in turn requires the intermediate processing, and timeliness of this process to be transparent, i.e. any intermediate processing used needs to be clearly identified, repeatable and automated on a predictable cycle, with appropriate exception reporting.

Roles of Point of Truth and Data Warehouse implementations

The above requirements for quality of provision do not dictate the component configuration, or the nature of organizations in the chain of provision, but they do demand certain implementation approaches.

Data Warehouse approaches provide a convenient single access point. These have the following advantages:

- can be used to force data provision into a common set of data access agreements;
- forces support for a small set of access protocols amongst consumers;
- can be used to promote common structures for similar data;
- concentration of effort/resourcing in robust engineering of solution to support high availability; and
- centralised monitoring of data access.

The data warehouse approach also has some inherent challenges and important risks. Those challenges that remain regardless of implementation quality include:

- Data Warehouse providers must deal with brokering and managing data access agreements, introducing a third party into all access arrangements;
- There is an continual need for resources to transform and load data from the various sources into warehouse repository; and

- There is a continual need to report data availability to end users.

The key risks (with inherent, but avoidable, potential for inappropriate implementation) include:

- Technology-platform-centric implementation, forcing all consumers to use certain technologies to access the warehouse;
- Upgrade: technology or data structure upgrades at warehouses change behaviour of services, forcing clients to adapt;
- Ad-hoc data manipulation: where the data manipulation to load data into the warehouse is a series of one-off exercises or based on a fragile chain of technologies, personnel availability and knowledge;
- Single point of failure: all users become reliant on a single organisation. Failure (technical or otherwise) of this organisation affects all members of the community; and
- Highly rigid structures, difficult to evolve.

Data warehouses can be safely implemented. However, this is only possible if the data used to populate these warehouses are available from stable sources and the ongoing existence of the warehouse can be assured both technically and financially. At this point data warehouses can provide the necessary transparency so long as they are kept up to date.

One opportunity then is that this updating process be undertaken automatically through an existing Point of Truth such that the data warehouse becomes a key client.

4.3.3. Scalability

Longevity and wide accessibility of the WRON require scalability. From a user-centric view, this has two possible interpretations:

1. Each new component added to the system causes no degradation of performance, usability or manageability of the existing system.
2. Each new component type should be designed to make it possible to trivially add many instances without requiring understanding or testing of impact.

The WRON must be scalable on a number of dimensions:

- Data Instances – amount of actual data records;
- Data Provision Services – number of distinct service provision points;
- Data Products – number of data types;
- Functionality of value-adding services; and
- Integration with external frameworks.

This leads to two architectural requirements – Service Profiles and Subscribe, not Describe which are described below.

Service Profiles

In a distributed environment, there are a number of implied contracts between users and service providers. These contracts are realised with software tools. From a scalability perspective, the set of types of contracts must be as small and coherent as possible so as to ensure the practicality of implementation.

For the WRON, these key contract types are published as service profiles. The WRON-RM will describe how service profiles are implemented and exploited.

Subscribe, not describe

There are a number of unscalable behaviours associated with the need to describe all artefacts within a distributed environment. These include:

- 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' (see Section 4.3.5) when recreating all aspects of metadata without re-use; and
- Impossibility of creating software components able to interpret and consume services based on ad-hoc descriptions.

The principle of 'Subscribe, not describe' encourages service providers to implement, where possible, services that adhere to published descriptions rather than fully describing each service instance. 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 (e.g. a service profile) than to describe it fully.

4.3.4. Simplicity

It is important that the WRON is as understandable as possible to its audience, and is simple enough to manage.

This does not however imply that the implementation or design can be over-simplified. Instead, the focus will be on dividing the WRON-RM into a set of viewpoints that address issues from the perspective of each stakeholder, so that the WRON is as simple as possible for each stakeholder, even if this means more complexity as a whole.

This strategy is common and effective – to create a series of modules that effectively interact. The cleanness of the modularisation is usually a function of separation of concerns:

4.3.5. Adequate Description

Descriptions (metadata) of component behaviour must be adequate for the intended use.

Resources should not be published until they are adequately described, and this means that in general, the type (behavioural metadata) and content domain (content, as exposed) need to be pre-registered descriptions before an instance of a service can be registered, even if the service is a unique component.

This implies that the developer of service profiles must take the end-use of the service as the driving factor. For example, a query that requires the End User to nominate a time period as a calendar month must fully specify the acceptable formats and range. A service that describes only the input and output data type is unlikely to be adequately described in terms of either discovery or exploitation of the content domain it exposes through those data types.

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

- Service type;
- Content exposed;
- Queries/invocation messages supported;
- Quality of Service; and
- Documentation (metadata) requirements.

A service profile's primary role is to ensure that the expected behaviour of a service meets the expectations of the service consumer, however it also is critical to:

- Document service behaviours in a reusable fashion;
- Provide service provider with an implementation checklist;
- Support automated conformance testing;
- Allow system managers to assess compliance; and
- Create similarity between different types of services to minimise overall system complexity (from both implementation and usage perspectives).

4.3.6. Who pays the cost?

The less effort expended for a given outcome the more cost-effective the solution. Thus the key factor in cost-effectiveness is to keep the effort required within an acceptable threshold for all stakeholders.

Table 4-2 gives an indication of the level of frequency different stakeholders will interact with the WRON. It can be seen that the greatest interaction will be undertaken by end users through Discovery or Access actions.

Table 4-2: Stakeholder actions and their frequency when interacting with the WRON

Stakeholder	Action	Frequency of activity
End User	Discover	******
End User	Access	*******
End User	Request functionality/data	***
Data Custodian	Deploy Service	***
Data Custodian	Describe Service	**
Data Product Designer	Describe Service	**
System Manager	Add Custodian	**
System Manager	Integrate domain	*

***** indicates expected frequency orders-of-magnitude i.e. ** = 10 *

It can be seen from Table 4-2 that the act of describing a product or service is significantly less frequent than end-user discovery and access. Therefore, it is worth investing in making discovery and use efficient, even if this means making the data description process more robust.

The WRON-RM will therefore aim to create the ideal architecture for the end-user needs, and to identify how to best support Data Custodians to meet these expectations.

4.3.7. Semantic Robustness

Interoperability ultimately derives from a need for a common understanding (Elkin 2007). That is, achieving interoperability requires an agreed point of commonality between stakeholders for certain practical purposes. For the WRON, this means that stakeholders must have the ability to declare their adherence to a common understanding, and by implication the ability to publish (or import) the necessary definitions. Other stakeholders can then choose to rely on these definitions.

Achieving this level of semantic robustness requires a means of managing definitions. This necessitates the creation of new, or engagement with existing, governance frameworks responsible for managing particular vocabularies.

It is important to note that semantic definitions should only be provided to the level of detail for which the governing body is comfortable.

4.3.8. No Private Contracts

No contract can exist between any two WRON services where each of the semantic components necessary for that contract cannot be traced back to an appropriate governance regime. That is, there must be no component of the WRON that depends upon unpublished private agreements which contradict or compromise the published standard. Such agreements directly threaten the long term viability of the WRON as they set a precedent for such agreements to proliferate.

4.4. Governance

The WRON aims to provide data access and information generation services to support the operational, planning and research requirements of a range of stakeholders. These stakeholders need to “know what they are getting”, in the same way that data and service providers need to “know what is being asked of them”.

Furthermore, the WRON must evolve to meet priorities, changing technical environments and data and tool availability. This evolution must be enabled without creating an unmanageable set of issues for the individual stakeholders.

If the WRON-RM is characterised as a set of policies that allow components to be safely added, these policies need to be clearly articulated and managed. Thus, appropriate governance at various levels is of vital importance to the ongoing success of the WRON.

4.4.1. Key Governance Regimes

Reference Model

The WRON-RM must be dynamic. That is, as new concepts and technologies become available, it is essential that the WRON-RM be able to address these and where appropriate, integrate them. This means that the WRON-RM must undergo a process of constant review similar that undertaken for many international standards.

Ensuring these reviews are undertaken suggests that the WRON-RM must be owned.

This governance regime will include:

- Ownership;
- Publication of revisions;
- Handling change requests;
- Assessing impact of changes and managing those changes;
- Adoption, Deprecation and Decommissioning of implementations; and
- Alignment of WRON-RM with external environments.

Version 0.1: The governance model for the WRON-RM is yet to be determined. This will be addressed in Version 1.0

Compliance Profiles

Compliance Profiles are the subset of Service Profiles for which tests can and should be implemented. As such, they will be the key artefacts that record the various types of service behaviours supported by the WRON. These behaviours can be viewed, in some respects, as types of agreements between providers and consumers.

Governance arrangements will need to be developed to ensure that such agreements can be developed in collaboration with both stakeholder types, and that the WRON itself acts as an efficient enabler of the publication and implementation of conformance profiles.

Governance may include provision for reference implementations, configuration examples, and compliance tests.

Vocabularies

Vocabularies are particularly important areas of concern for governance (Renaud 2004). They are often highly dynamic and provide the main point of interaction between members of a community when negotiating interoperability. The WRON must support the ability of vocabularies to evolve over time.

Registries and Registers

A key interaction for many stakeholders in WRON will be Discovery. This includes discovery of:

- services that meet certain requirements;
- service profiles; and
- definitions.

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

Each register will need a formal governance regime (c.f. ISO 19135 Geographic information) determining who has rights of access.

Each registry will need separate governance to establish individual registers and delegate governance to appropriate owners. This will also apply to semantic registries and their contents, e.g. domain models, service instances and meta-information.

Data Resources

Individual data resources may be relied on by users and providers of intermediate (processing or ‘value-added’) services. Therefore, the governance of resource availability is important.

Specific requirements include:

- Regular performance/availability monitoring;

- Conformance/compliance monitoring;
- Usage monitoring; and
- Failure handling.

4.5. WRON 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. The WRON-RM 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.

In this section, the architecture is described according to the perspectives of different types of stakeholders and necessary management roles. 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 key perspectives that have been identified for this section are:

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

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

Each perspective is organised to make it easy to understand the WRON system responsibilities. The goal is to ensure each perspective is as simple as possible without compromising the ability of the WRON to be progressively implemented in an effective, consistent manner.

A common framework is used to describe each perspective. Table 4-3 shows the information included in this framework.

Table 4-3: Information Sections for WRON System Use-Case Perspectives

Section	Description
Use Cases	The Use Cases section describes the operation of the WRON, from the appropriate perspective, in detail. The Enterprise Viewpoint provides a “what needs to be achieved in the end” overview. Different implementation patterns will be described in the Computational and Engineering Viewpoints.
Key Roles	This section takes the form of a diagram showing the entry points into the WRON from the stakeholder’s perspective, and the key Actors (Roles) that the stakeholder will have explicit or implicit knowledge of, and whose delivery of services via the WRON will form part of the business function.
Integration Points	The Integration Points section describes how the WRON can be exploited through the stakeholder’s existing or future operational environment. This is a key element that the WRON needs to deliver to meet its business objectives – an infrastructure providing components seamlessly integrated into stakeholder activities.
Key Artefacts	This section identifies the artefacts that must exist for the different stakeholders to fulfil the contract they require. A contract is a formalised agreement between parties (e.g. data providers and the WRON) that

	defines the obligations of and benefits to both sides. These contracts provide the link to the “how will the WRON work” details. Diagrams identifying (but not describing) the artefacts, in the simplest possible fashion, are provided in this section with full descriptions of the components provided in the relevant viewpoint (Information or Computational). The intent is to provide a complete list of all contract types that must be enabled by the WRON, and to keep these to as small a set as possible.

4.5.1. End User Perspective

The End User is a person (or organisation) who is using the WRON to obtain information. The End User may interact with the WRON either through tools that directly implement WRON-compliant interface(s) or indirectly through accessing information that has been delivered via the WRON. Examples of end users include water regulators, water supply companies, irrigation companies, catchment management groups, government departments and the public.

Details of how the End User interacts with the WRON and the people or systems involved in those interactions are described in the following sections and illustrated in Figure 4-2: Use Cases and Participating Roles from the End User’s Perspective

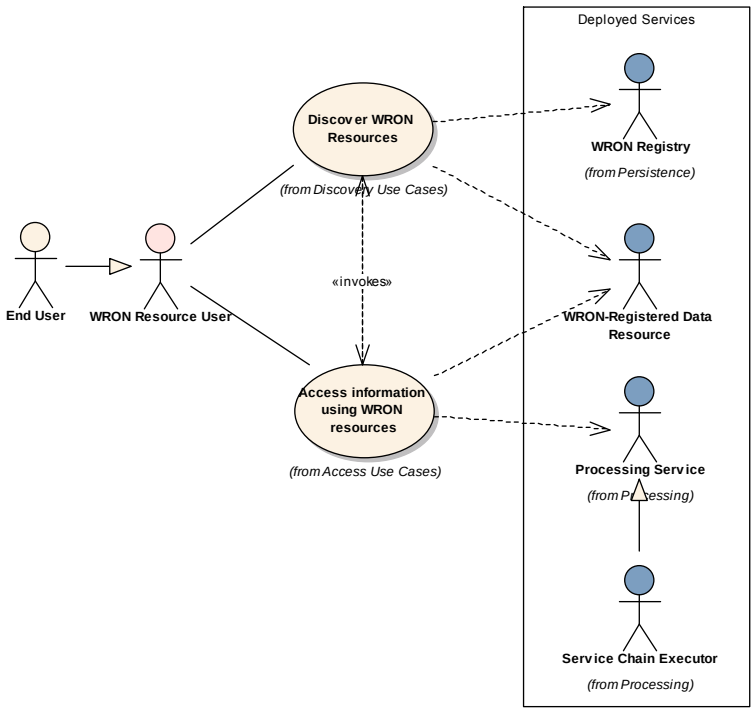

Figure 4-2: Use Cases and Participating Roles from the End User’s Perspective

Dam Level Scenario: End Users

User: Bob, Analyst, Three Rivers Basin Commission

Bob is an analyst at the Three Rivers Basin Commission. He is responsible for producing monthly water account reports for the commission. As the commission is a multi-jurisdictional institution, information about current dam levels (stocks) at different locations within the Three Rivers Basin is disbursed and must be gathered and collated to produce the water account.

User: John, Finance Manager, ABC Rural Bank

John is a Finance Manager at the ABC Rural Bank. The bank mainly focuses on rural customers and must stay current on conditions affecting the agricultural industry. When John evaluates the risk involved in a loan application, climatic conditions and trends need to be included as they form part of an economic indicator for the industry. John would like to know the trends in dam levels for his region.

Use Cases

The End User's perspective consists of two clear use cases:

1. Discovery; and
2. Access.

Discovery Use Case

The Discovery Use Case concerns the ability of an End User to discover resources or capabilities, to meet a particular need. In general, the End User will interact with the WRON through a WRON-compliant user interface, either directly, or through an intermediary.

The resources and infrastructure involved in providing access to data and services will be largely transparent to the End User. However, the End User should be aware of the existence of the WRON Registry and data or value-added services and their role in facilitating access to raw or processed information.

Discovery functions may be embedded in data access workflows, where interoperable data resources are automatically selected according to the execution context (such as regional data services for a common model). There will also be discovery functions embedded in other design and deployment use cases, to discover the information required to describe a new or subscribe to an existing service profile or component design.

A key principle of the Discovery use case is the concept of adequate description (see Section 4.3.5) which seeks to ensure that additional burden is not imposed on the End User (or data/service provider) in order to find and use data or services within the WRON framework.

The WRON-RM does not assume any particular discovery paradigm (portal, desktop, etc). It is simply concerned with enabling the process.

Access Use Case

The Access use case concerns the ability of the End User to access information using WRON resources or capabilities.

From the End User's perspective, data access is fundamentally a query operation. This query will usually be defined by the End User through an interface (e.g. selecting a region on a map or selecting options from a list). Underlying the interface are query models. These provide mappings from data access services to query definitions (which may be parameter lists, filters or SQL statements). Query models may be published in a registry and attached to service instances or profiles. Semantic information and reasoning may then be exploited to determine a mapping between the End User's query and available query models.

Exploitation of a query model may result in the direct invocation of a data access service. However, it will more probably result in the composition of a service chain to return the requested information product. A service chain is an arbitrarily long composition of services that delivers a specific information product. Service chains may be published for specific products, but can also be composed dynamically through composition services using semantic reasoning.

Dam Level Scenario: Discovery and Access Use Cases

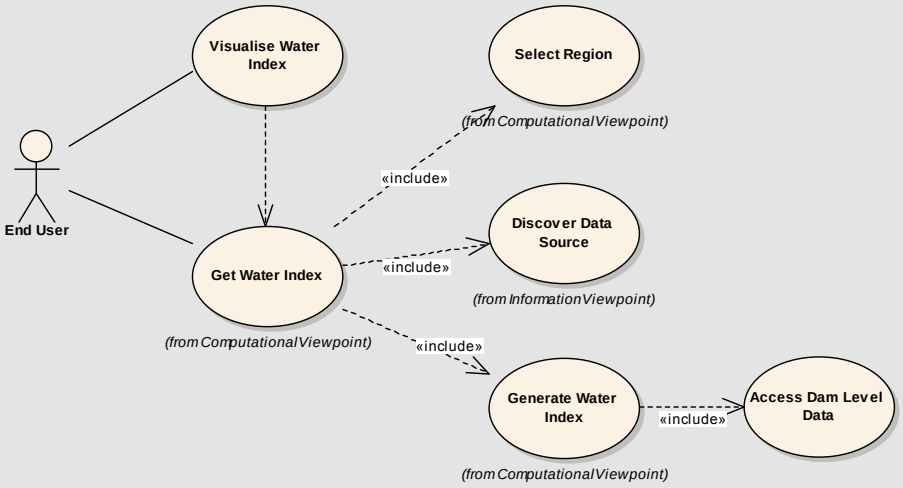

Bob needs to obtain the Water Index (the measure of current storage levels compared with capacity) for the dams in his area of interest so he can compile a monthly report. He uses an application called WIRT (Water Index Reporting Tool) which is a WRON Compliant tool. He first selects the region containing the dams he requires water indices for. WIRT then requests a list of data sources from the WRON registry that provide dam level data (Discovery). This information is supplied to WIRT which then determines the most appropriate source to access. A request is then made to this source (Access) and the information required by Bob is delivered.

Key Roles

The key roles from the End User's perspective (the people and/or systems the End User may interact with – explicitly or implicitly) are:

- WRON Resource User;
- WRON Registry;
- WRON-Registered Data Resource; and
- Processing Service

An End User is a type of WRON Resource User. WRON Resource Users are entities (human or machine) whose primary interest in the WRON is as a means for accessing information. WRON Resource Users may also play the role of Functionality Provider (see Section 4.5.3). An End User does not play this role. Whether a regulatory authority, an irrigator or a householder concerned with current dam levels, the End User is only concerned with finding and exploiting the information available within the system.

A WRON Registry exists to support discovery by allowing identification of available resources. The registry must be capable of supporting functions that identify what resources are available through processing of other data. This function may also be provided by a specialised processing service – hence access may be called iteratively by various stages of the discovery

process. For example, discovery may access WRON resources to provide preview functionality. Conversely, access may invoke discovery to choose a set of services to access at run-time.

A WRON-Registered Data Resource provides data. The End User interacts with data resources through the Access use case where data is provided in response to a query.

Data may require some form of processing before delivery to the End User. This is the role of Processing Services. To the End User, some Processing Services may appear as data access services with the processing occurring behind the interface. As an example, a Processing Service may provide access to data that is translated before delivery. The End User only sees the interface to the translated data, and so the translation service appears as a data access service. In other cases, the End User may deliberately select particular Processing Services (e.g. a model) to transform data accessed directly.

Dam Level Scenario: Key Roles

In Bob's case (described above), he has interacted with the WRON Registry – to discover sources of the Water Index Information as well as Processing Service - that provided the Water Index. This Processing Service has, in turn, interacted with a WRON Registered Data Provider – providing the required Dam Level data.

Integration Points

WRON facilitates the integration of WRON data and services with stakeholder's existing or future operating environments. The End User should interact with the WRON through WRON-compliant interfaces or services. It is expected that WRON Compliance, from a client software point of view, will be the ability to interoperate with one or more WRON service profiles, which may in fact be designed specifically to enable integration with common technology platforms.

The WRON's purpose is to maximise accessibility of data through well-known service profiles. For systems that are not WRON compliant (e.g. legacy systems), WRON wrappers can be developed to automate integration. The wrapper might perform content transformation or semantic translation to adapt the non-compliant system into a compliant one.

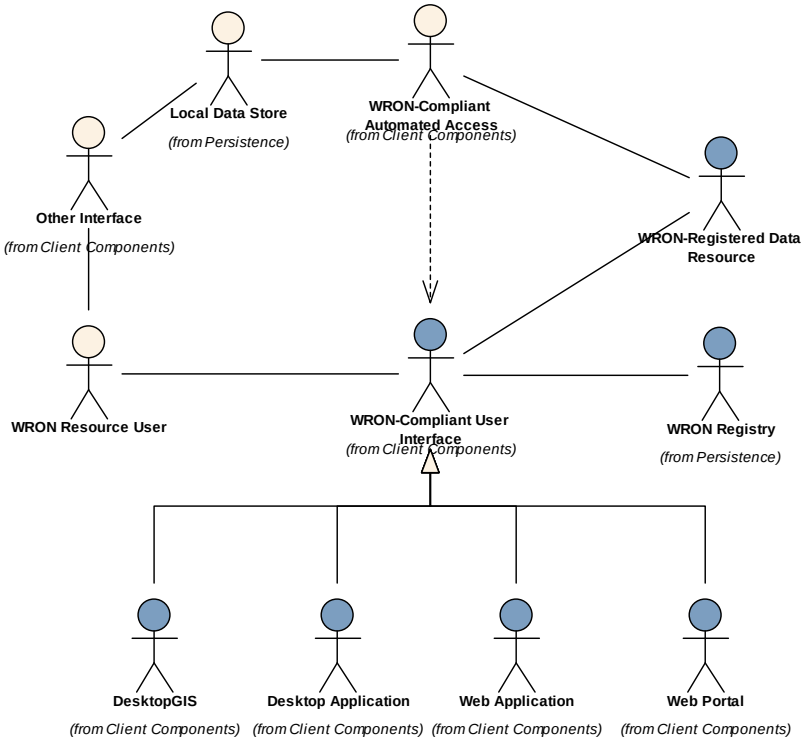

Figure 4-3: System Integration Roles

Figure 4-3 highlights the requirements for this set of WRON capabilities to be integrated with the stakeholder's existing or future operating environments. It illustrates how the WRON Resource User (i.e. the End User) can interact with the WRON through a number of applications and interfaces. The user may interact directly with a WRON-compliant interface to new or legacy systems such as a desktop GIS application, web application or portal. Alternatively, they may interact through some other interface that accesses WRON data behind the scenes and provides results through a local data store. WRON-compliant services and interfaces are required for access to WRON-registered data resources.

Dam Level Scenario: Integration Points

Bob integrates with the WRON via WIRT. WIRT has been specifically developed to be WRON Compliant and hence contains a WRON Compliant interface. This interface is capable of performing a WRON Registry request as well as supports queries and data associated with Water Indices. This interface does not, however, support other WRON Services Profiles.

Key Artefacts

Figure 4-4 shows the key information exchanges that underlie the Discovery use case from the End User's perspective. The WRON is particularly concerned with the responsibilities of external stakeholders in enabling End Users to discover and access relevant information in a timely fashion. End Users rely on a small, coherent and consistently described set of service profiles to locate and understand conformant services. As such data resources must conform to published service profiles to allow for discovery and use. By subscribing to an existing profile, individual data providers reduce their burden of description. Service Profiles must be available via the registry to allow specialisation of coherent families of similar services.

In some cases an End User will play the role of a Service Chain Composer. This will occur when an End User manually composes a service chain. In other cases an End User may interact with a Service Chain Composer to achieve a required outcome.

A Service Chain Composer exploits the registry (which has knowledge of what data and/or services are available) and the Service Chain Executor in order to provide an appropriate service chain. The Service Chain Composer holds knowledge of what product is required, and is able to determine the capabilities of available services. The interface between the End User and the Service Chain Composer may allow interactive composition of the service chain, or may present a set of virtual data sources that encapsulate published service chains. The service chain is then executed by the Service Chain Executor.

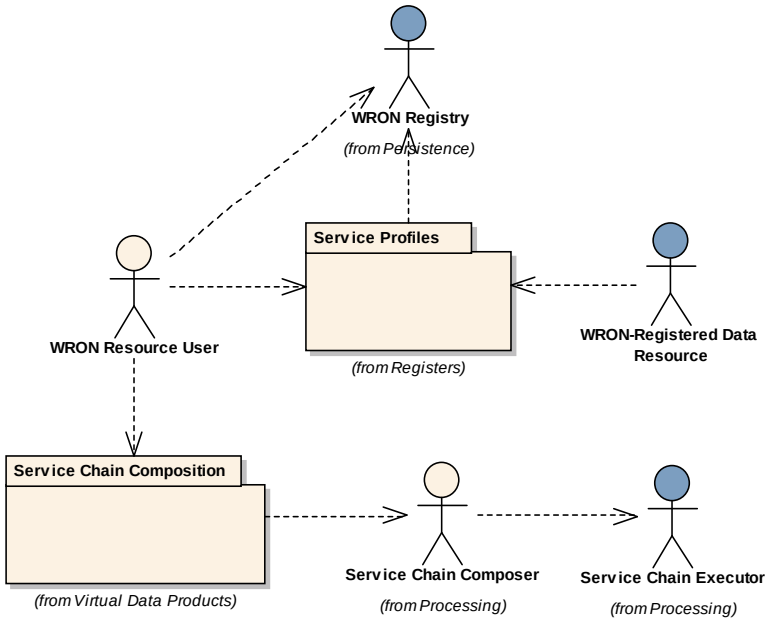

Figure 4-4: Key Contracts and Artefacts in the Discovery Use Case

Note that Figure 4-4: Key Contracts and Artefacts in the Discovery Use Case

does not highlight the possible components (such as a discovery client), but the fact that the End User understands that a WRON Registry can be accessed to discover resources. This registry may be discovered in many ways, or hidden behind a familiar access point.

The End User interacts with the WRON via a WRON-Compliant User Interface. This interface relies on the existence of Query Models and Service Profiles for access to WRON-registered data resources. This as shown in Figure 4-5.

Query Models allow user defined queries (Query Instances) to be mapped to appropriate data resources by matching the query against the resource's service profile. Thus a query model defines the queries that are supported by particular service profiles and hence provides definition for possible 'contracts' between Resource Users and Data Providers.

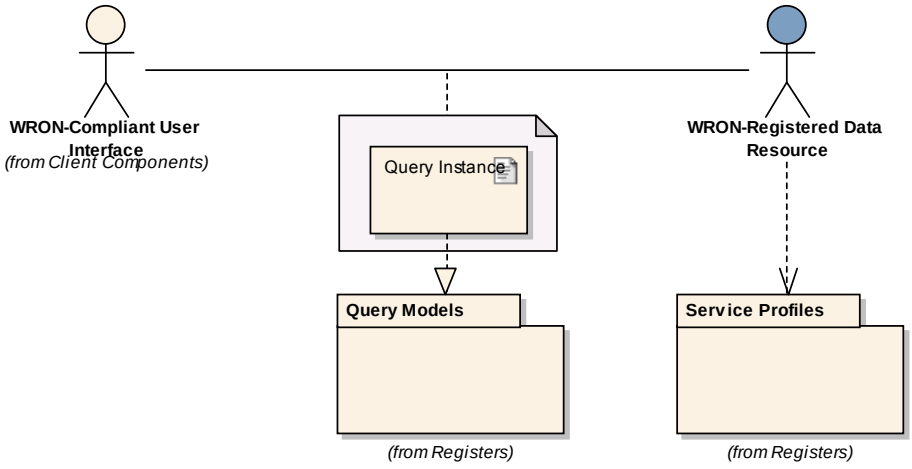

Figure 4-5: Key Contracts and Artefacts in the Access Use Case

Dam Level Scenario: Key Artefacts

Version 0.1: Text describing the key artefacts in the context of the Dam Level Scenario will appear here in WRON-RM 1.0

4.5.2. Data Provision Perspective

Data Providers are individuals or organisations that provide access to data through the implementation of WRON-compliant data services. Examples of possible data providers include water reservoir managers, irrigators, State and Federal Government departments and research organisations such as CSIRO or NASA (remote sensing data).

A Data Service Provider is a type of WRON Service Provider and may not be the actual custodian of the data. This relationship is shown in Figure 4-6 along with other key roles and use cases involved in data provision. Details of these roles and use cases are provided in the following sections.

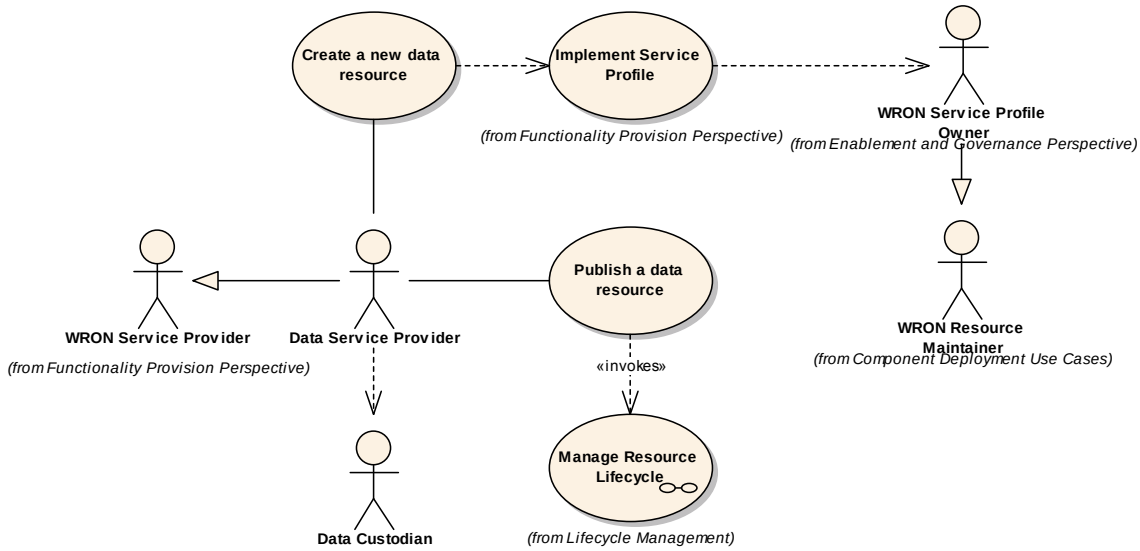

Figure 4-6: Key Use Cases and Roles from the Data Provider's Perspective

Dam Level Scenario: Data Providers

Provider: Trish, Engineer, Three Rivers Dam

As an engineer at the Three Rivers Dam, Trish is responsible for managing dam level data which she provides daily to a range of interested parties. This data is delivered via a WRON Registered data service. This delivery method frees Trish up from answering ad-hoc queries throughout her already busy day.

Use Cases

Key use cases from the Data Provider's perspective are shown in Figure 4-6. Primarily, a Data Provider is concerned with four use cases:

- Create a new data resource;
- Implement service profile;

- Publish a data resource; and
- Manage resource lifecycle.

Create a New Data Resource Use Case

The WRON-RM assumes integration of data through web-accessible services, either via live (on-demand) access or automated distribution of updates. Each data provider will undertake delivery of services that conform to a WRON-registered service profile. The WRON will not dictate the internal architecture of each service, only its behaviour to clients.

This use case encompasses the internal processes that a Data Service Provider will complete in order to produce a data service. These processes include the business aspects of providing access to data such as allowing access to the database, ensuring the data is accessible through firewalls, and arranging for maintenance of the data resource. As these processes are internal, they will be different for each data provider.

It is expected that useful best-practice organisational procedures can be developed and disseminated, and the technical aspects alone are not sufficient to meet WRON requirements, hence this Use Case is part of the WRON-RM.

The technical aspects of creating a WRON data resource by implementing a service profile are described in the 'Implement Service Profile' use case.

Implement Service Profile Use Case

A Data Service Provider creates a WRON data resource by implementing a service conforming to a published service profile. This profile may already exist, or may be created by the data provider. Conforming to an existing service profile reduces the effort required to describe the data service (c.f. the architectural principle of Subscribe, not Describe). Where an existing service profile does not specify sufficient detail, the data provider will classify the behaviour and content of the data service by extending the most relevant common profile with a new specialisation.

Implementation of a service includes conformance testing of the service against the profile to ensure compliance. Following conformance testing, the new data resource is published, thus allowing WRON resource users to discover it. The conformance test reduces the effort for data publication by providing a quick and easy means of testing much of the service implementation.

Publish a Data Resource Use Case

Data resources require publishing to enable discovery and access. Publication involves describing the new resource, and registering it with the appropriate registry. The description effort required should be minimal as the service profile will provide much of the required description. Publishing a data service also initiates a lifecycle management process which ensures the resource is properly maintained while it is part of the WRON. Part of this management process is notification of changes to the resource.

Commonality of lifecycle management is necessary to make the WRON cost-effective to participate in, since tools can be developed to support these Use Cases. It is also necessary to minimise user uncertainty over service availability.

Manage Resource Lifecycle Use Case

The WRON will provide a common lifecycle management framework for all resources, so that it can be predicted exactly what will be available for each resource, and how each resource will behave.

The WRON resource lifecycle management processes form part of the agreement between individual Data Service Providers and the WRON. As such, they allow the WRON to manage this agreement with an arbitrary number of stakeholders, in a consistent fashion. (Each stakeholder is likely to have interests in many components of the WRON, and a common regime is critical for this to be managed).

Dam Level Scenario: Data Provision Use cases

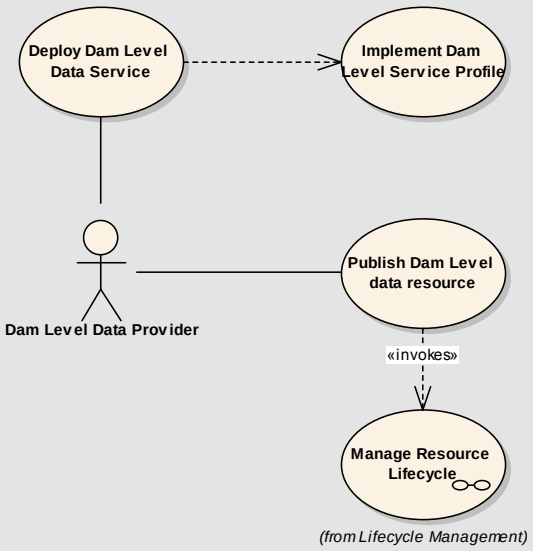

Trish engages a technical team to deliver a new WRON data resource. Since Dam Level data is commonly required, a service profile for "Dam Levels" already exists. The existing service profile provides an unambiguous specification of exactly what needs to be provided making the process simpler for all concerned – it provides Trish with the business outcomes she is required to meet, and technical detail for her technical team.

Trish then publishes the Three Rivers Dam Level data resource in the registry and signs up for notifications in case of changes. Her existing clients requiring Dam Level data can now discover and access the data without having to contact her directly.

Key Roles

The key roles in the Data Provision Perspective, as shown in Figure 4-6, are:

- Data Service Provider;
- Data Custodian; and
- WRON Service Profile Owner.

A Data Service Provider is a type of WRON Service Provider. Data Service Providers provide services that allow access to data resources.

Data Custodians are individuals or organisations that own or maintain data resources. Data Service Providers may be Data Custodians, but this is not a requirement. That is, the provider could be providing the service on behalf of a custodian. In this case, an agreement should exist between the Data Custodian and the Data Provider covering what and how the data is to be published.

Data Service Providers implement service profiles. These profiles are owned by WRON Service Profile Owners – a type of WRON Resource Maintainer. The service profile thus provides an additional function of providing a specification for the technical aspects of any required service establishment contract between a custodian and a service provider.

Integration Points

Service Configuration

One advantage of the use of common patterns for machine-readable service profiles is that it becomes possible to use these to assist in configuration of data provision services. A large

proportion of the configuration is description metadata, and this can be automatically derived from the service profile, or turned into a series of simple choices from controlled vocabularies (e.g. using a Wizard style process).

Service Registration

Web service interfaces to support conformance testing and registration of new services will allow integration of service registration with the provision of the services themselves.

Dam Level Scenario: Integration Points

Version 0.1: Text describing the integration points from the Data Provider perspective in the context of the Dam Level Scenario will appear here in WRON-RM 1.0

Key Artefacts

A data provider delivers a service that is immediately useful within the WRON environment. There should be little or no additional burden for mediation by the data provider with each new data access.

To achieve this, the agreement between the data provider and a resource user is mediated by the WRON in two ways:

- The terms of use; and
- The technical specification (i.e. the service profile).

The WRON-RM defines a small set of terms of use for standard agreements that the data providers can subscribe to. This allows the WRON to mediate access according to its common governance rules - i.e. the WRON will support a small set of simple user roles that can be easily and predictably matched with individual data services.

This allows for the majority of cases to be handled simply. It also allows for data providers to invest in creating and propagating into the WRON any additional approaches they feel are necessary. The burden of effort is thus on the stakeholder who wishes to add complexity to the system, not the end-user who needs the simplest possible environment to meet their needs.

The technical details of the agreements must also be encapsulated. This is achieved through the Service Profile which binds all the various aspects together, and allows a Data Service Provider to simply implement the Service Profile and use it to test conformance. The Service Profile can also be used to populate the majority of metadata required for the service to be adequately described for discovery and use.

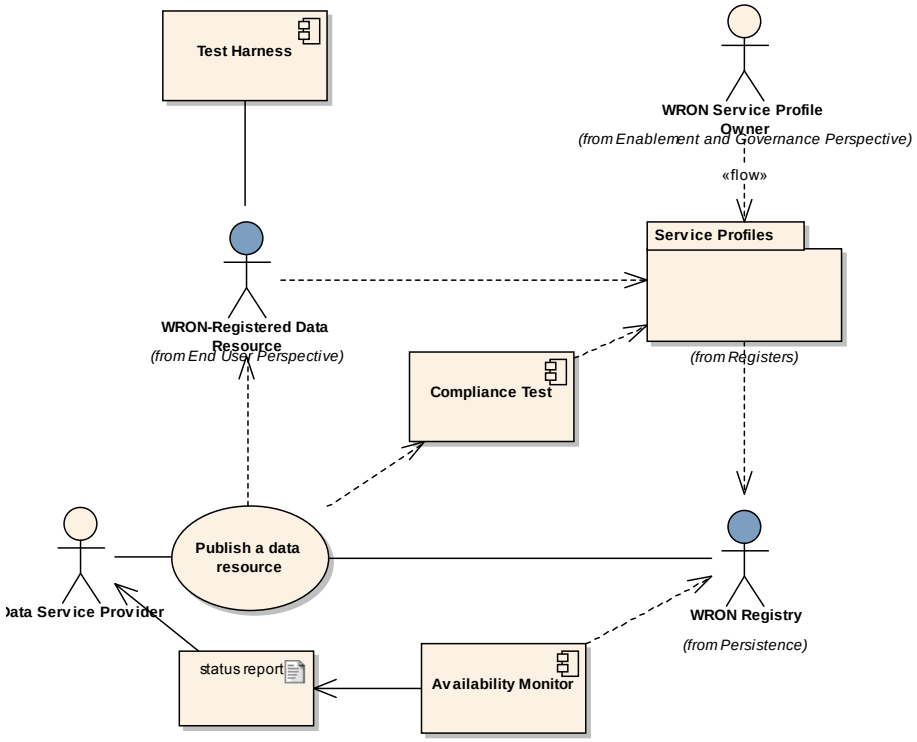

Figure 4-7: Data Provision Key Artefacts

Figure 4-7 illustrates the components and artefacts involved in publishing a new data resource. The publication process for a resource includes testing compliance with the service profile. Part of lifecycle management of the resource is availability monitoring and notification as indicated in the diagram by the status report.

Dam Level Scenario: Key Artefacts

Version 0.1: Text describing the artefacts from the Data Provider perspective in the context of the Dam Level Scenario will appear here in WRON-RM 1.0

4.5.3. Functionality Provision Perspective

Functionality Providers provide processing or value-add services to the WRON. These services take data from a data service and transform it in some way. Examples of processing service tasks include data translation, visualisation, modelling and prediction. Providers of these services might include businesses with a requirement to report on water availability or quality, or research organisations such as CSIRO and CRCs.

Figure 4-8 shows the key roles and use cases in the Functionality Provision perspective. These roles and use cases are detailed in the following sections.

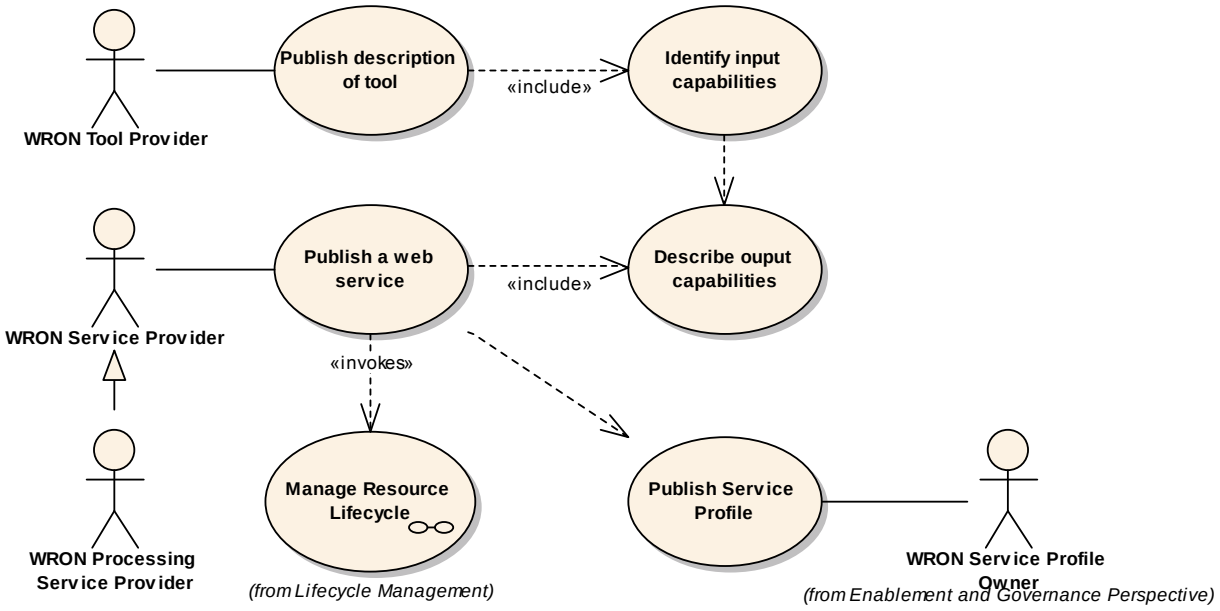

Figure 4-8: Key roles in the Functionality Provision perspective

Dam Level Scenario: Functionality Provider

Provider: Smith Water Corporation

Smith Water Corporation is a bulk water delivery business responsible for water delivery to town water supplies, industry, irrigation, stock and environmental flows. It owns 21 dams and manages another 11 dams on behalf of the Government. Smith Water supplies “water index” reports to the State Government so that available water determinations and allocations can be made.

Use Cases

Figure 4-8 identifies the key use cases from the Functionality Provision perspective. These include:

- Publish description of a tool; and
- Publish a web service.

Publish Description of a Tool Use Case

Tools are software components that may be deployed by service providers.

A tool is a distributable component designed to implement a specific function. It is not the actual function deployed as a service. Giving WRON stakeholders the ability to create interoperating functional components is critical. Hence the WRON infrastructure and the WRON-RM need to explicitly support tool description and deployment.

Tools can be seen as types rather than instances. They allow a growing suite of common WRON functional components to be encapsulated within an evolving environment as software implementations.

Examples of tools include an analytical model or a Web Feature Server implementation. As with services, tools need to be adequately described and published before they can be discovered and used. It is especially important that the input capabilities of tools are described, since this will determine the ability of the WRON to support delivery of data to users of the tool.

Publish a Web Service Use Case

As with tools and data resources, web services must be described and published to enable discovery and use. Description again follows the pattern of subscription to a published service profile. Lifecycle management is also initiated for published web services and the Processing Service Provider can register for notification of changes.

The use cases shown in Figure 4-8 provide a high level pattern for the delivery of processing services. While this basic pattern for functionality provision will be reusable, the common implementation platforms (i.e. frameworks or toolkits used for implementing processing services) used to implement the WRON initially will each have different procedures.

Specific use cases need to be developed for each class of functionality (e.g. models, data aggregation, data translation and portrayal services) to identify the particular constraints that apply for each platform. The goal of these use cases will be to make the WRON-RM immediately accessible to stakeholders familiar with a WRON compatible technology platform as a clear set of simple tasks.

Key Roles

The use cases for the Functionality Provider perspective describe how WRON Functionality Providers will deploy tools and services to let WRON Resource Users exploit resources available through the WRON. The WRON will provide the underlying data access framework for all distributed data access requirements. In general, it should be possible to integrate WRON resources into an external environment, or to make an environment accessible to the WRON user community.

From an engineering perspective, there are various ways of deploying functionality – as a reusable component (a Web/GRID service), a deployable configuration for an execution service (e.g. a configuration for a model execution service) or as a self-contained application (e.g. a downloadable widget or executable application). The two main categories of Functionality Providers considered here are Tool Providers and Processing Service Providers.

WRON Tool Provider

A Tool provider makes functionality available to a particular environment through the provision of reusable components. Examples of functionality delivery environments include:

- Desktop User Interface – distributable download (e.g. an application such as ESRI ArcExplorer);
- Desktop User Interface – application configuration (e.g. a project file that specifies where data resides);
- Standalone Application;
- Widget (tool loaded into framework);
- Web page; and
- Web Portlet (or other web based UI component).

WRON Processing Service Provider

A WRON Processing Service Provider is a type of WRON Service Provider who makes available processing services that provide a value-add to the service's inputs. Examples of processing services include simple averaging or trending services through to complex process models.

A key initial focus of the WRON will be Web accessible processing services. These allow encapsulation and re-use of key algorithms and models recognised by the WRON user community. The extensibility of the WRON is, however, critical, and the methodology for deploying, describing and exploiting such services requires significant development and refinement.

Dam Level Scenario: Water Index Service Provider Use Case

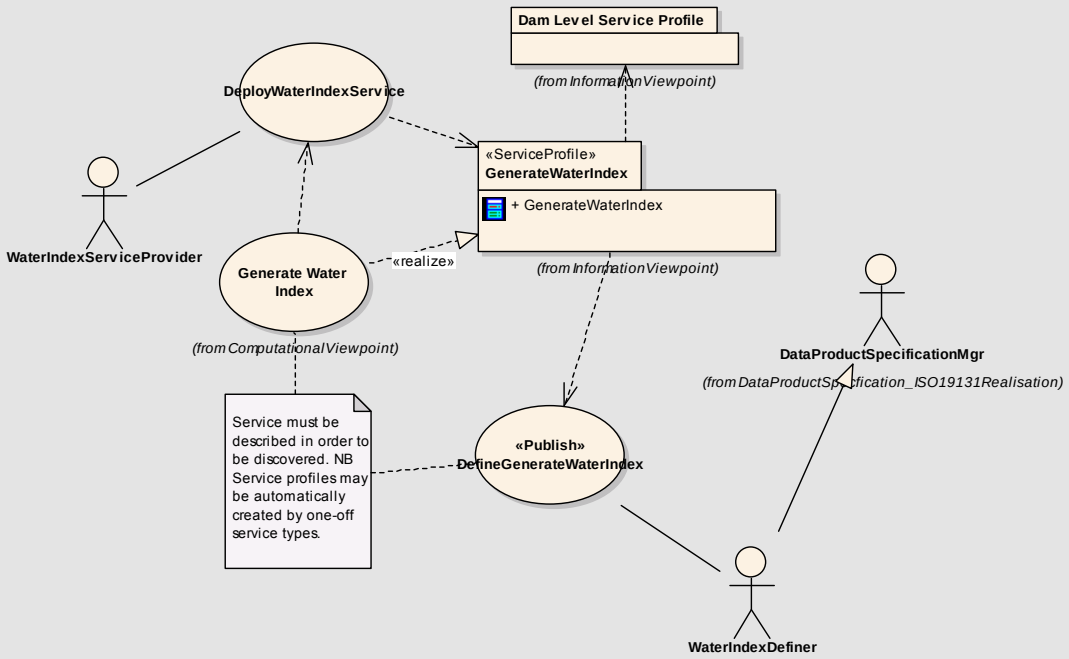

Smith Water provides the water index service based on a specification determined by the water index definer who might be the state government. The water index specification exists as a water index service profile.

Integration Points

The WRON will include components that support processing of data by reusable services, such as models, that may be invoked by third parties and deliver data or User Interfaces (UI) to WRON Resource Users. There will be a number of key technology platforms that will support access to such functionality.

This viewpoint is primarily focussed on the business requirements of the WRON, and how the WRON may adopt such platforms or component services as they become available. This must be achieved in a consistent fashion. Specific integration priorities will need to be identified in consultation with stakeholders. It is expected that WRON resources (data and services) will be made available to both WRON Resource Users and externally managed environments through a relatively small set of reusable interfaces such as those shown in Figure 4-9.

Figure 4-9 provides examples of possible interfaces between WRON tools, services and applications. The Load Model, Execute and Access Data interfaces allow interaction between tools such as the Service Chain and Model execution frameworks, data access and processing services, and desktop or web applications. The use of well defined interfaces reduces the burden on tool and service providers by providing specifications that can be implemented consistently.

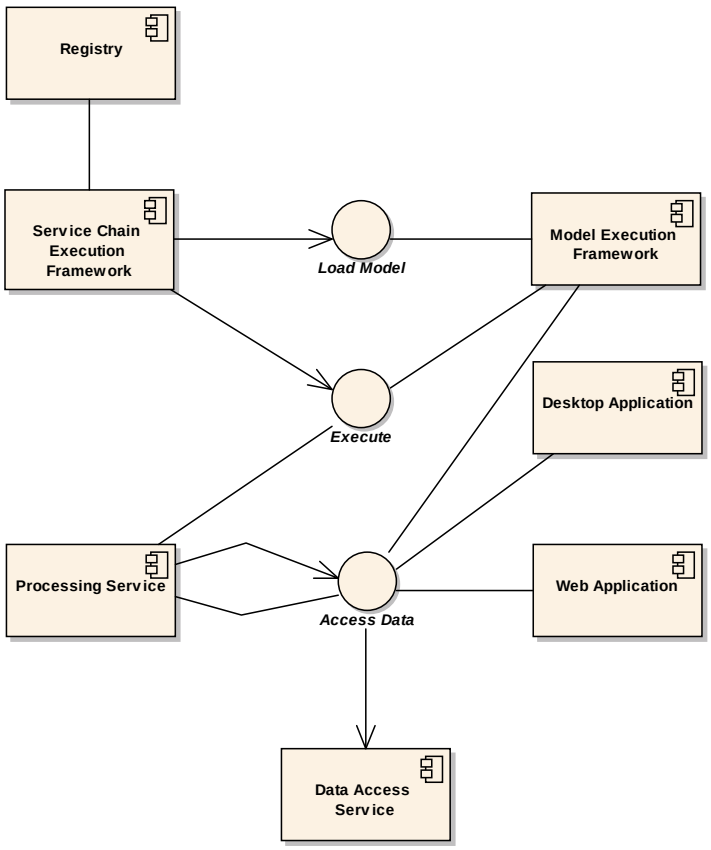

Figure 4-9: Examples of Interfaces Between WRON Resources and Resource Users

Dam Level Scenario: Integration Points

Version 0.1: Text describing the integration points from the Functionality Provider perspective in the context of the Dam Level Scenario will appear here in WRON-RM 1.0

Key Artefacts

The key artefacts from the Functionality Provision perspective are:

- Service Profiles; and
- Functional Component Types

Service Profiles

Tools and services are described by providers using service profiles. Existing service profiles for similar tools or services may be subscribed to and extended (by describing new features, or restricting functionality) to reduce the burden on the provider. For the purpose of ensuring interoperability it is necessary that service profiles for WRON functionality must ensure that the functionality described integrates with other WRON functionality.

Functional Component Types

The WRON provides a flexible framework that allows components to be delivered in a variety of ways, using any available technology. This is needed to ensure the WRON remains relevant in a climate of constantly changing technologies and environments.

The execution environment of any WRON value-add service must be clearly identified, with a goal to maximise availability. Exact bindings will be discussed in the Technological Viewpoint, however the following may be assumed:

- (Composable) Web services (eg data services, model execution services);
- Web Applications;
- Deployable client applications (deployable tool); and
- Integration into client applications.

Figure 4-10 illustrates the role of service profiles in describing service and tool capabilities. It also provides an example of how the system can be extended by the registration of new information artefacts.

A Modelling Framework Provider may offer a model execution service. The model itself is published to a Model Code register with a service profile describing the model capabilities, inputs and outputs to enable its use. The link between the new information artefact (the model code) and the existing artefact provides enough information for users to discover and access the modelling framework as a series of virtual services by binding the model code with the model execution service within a service profile. The resource user sees the data product offered as a result of the availability of this registered capability.

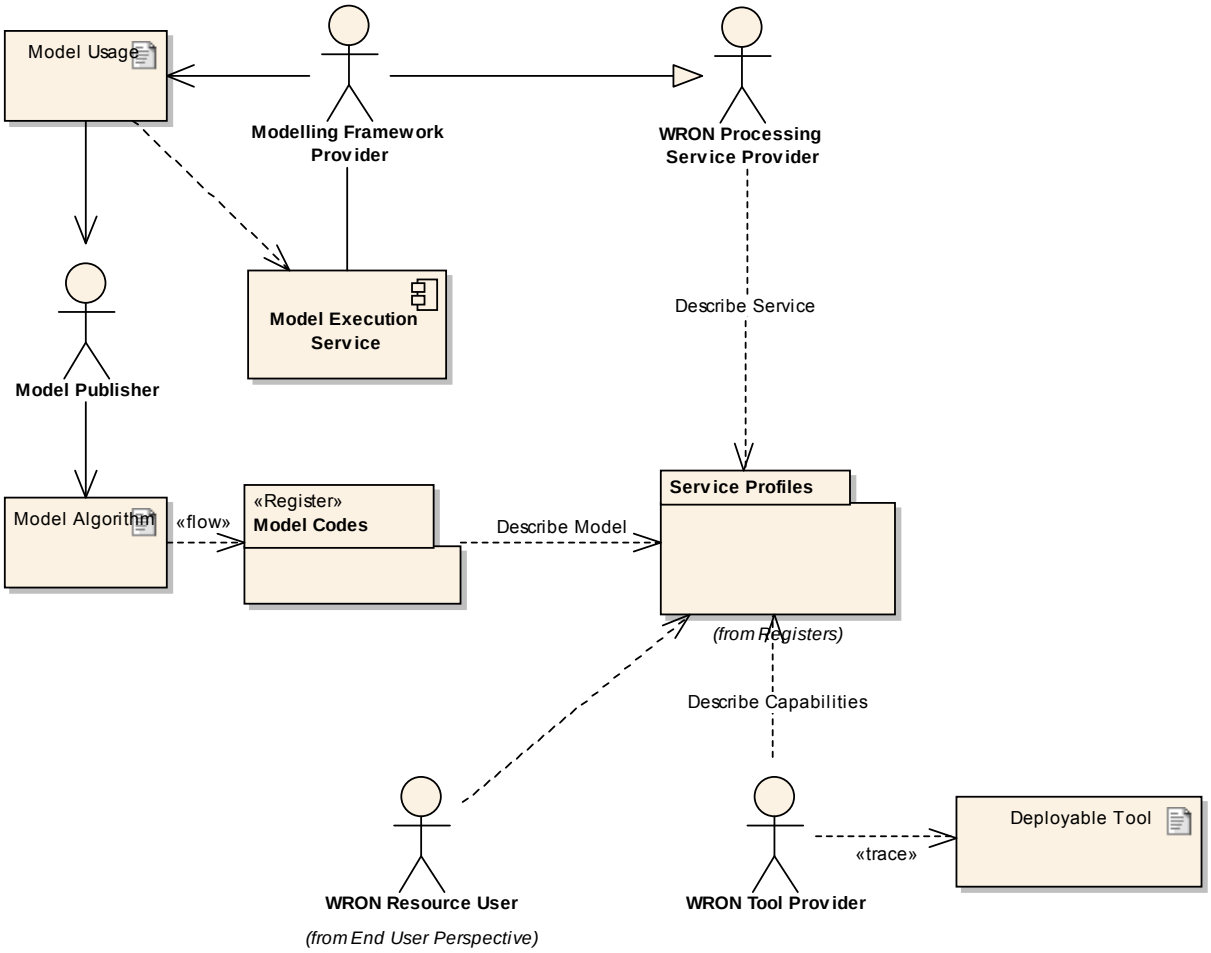

Figure 4-10: Information artefacts and contracts for functionality (processing service) provision

Dam Level Scenario: Key Artefacts

Version 0.1: Text describing the key artefacts from the Functionality Provider perspective in the context of the Dam Level Scenario will appear here in WRON-RM 1.0

4.5.4. Enablement and Governance Perspective

Version 0.1: The description of the Enablement and Governance Perspective here is not complete. It will be complete for WRON-RM 1.0

The usefulness and success of the WRON will depend fundamentally on its reliability. This reliability applies primarily to the behaviour of the resources that can be accessed through WRON and the predictability of their behaviour. If this is unconstrained, end users will not be able to survive the amount of variability and uncertainty that they encounter, and the level of effort required to deliver data resources will be unacceptably high.

4.5.5. Cross-Business Domain Integration Perspective

Version 0.1: The description of the Cross Business Domain Perspective here is not complete. It will be complete for WRON-RM 1.0

Domain Integration Use Cases

These Use Cases describe the functions within the WRON to enable integration of resources from external domains within the WRON.

4.5.6. System Maintenance Perspective

Version 0.1: The description of the System Maintenance Perspective here is not complete. It will be complete for WRON-RM 1.0

Given that the WRON is a long-lived, evolving and flexible framework with multiple stakeholders, system maintenance practices will need to be carefully developed. Provision must be made to ensure ongoing operation, monitoring of usage, performance and reliability of the systems involved. Further detail will be added to this section in Version 1.0.

5. Information Viewpoint – Key Artefacts

The Information Viewpoint of the WRON-RM elaborates the key information artefacts defined in the WRON RM-0.1. The WRON interoperability framework specifies the modelling approach for all categories of information the WRON deals with, including their thematic, spatial and temporal characteristics as well as their meta-information.

The WRON-RM does not specify an information system. Instead it provides a framework for distributed information systems and WRON Applications based on a service-oriented architecture. As such, the Information Viewpoint of the WRON-RM provides an integrated specification framework in order to support a formal specification of conceptual WRON-RM information meta-models and metadata-models in the context of WRON Applications.

The information viewpoint is broken up into two sections:

1. **Key Artefacts:** Identifies the nature, and extends the description, of key information artefacts defined in the Enterprise Viewpoint.
2. **Interoperability Framework:** Specifies methodologies for ensuring that the WRON as a system can grow and integrate to meet long term objectives.

5.1. Introduction

The Key Artefacts section takes the key information artefacts identified in the business-needs oriented Enterprise Viewpoint and further describes their composition and management.

The artefacts described below are cross-referenced against the same stakeholder perspectives used in the Enterprise Viewpoint, recognising that they may be relevant to multiple stakeholder groups.

The artefacts identified here are expected to be directly visible to stakeholders in order to support the Use Cases identified in the Enterprise Viewpoint. This set of Use Cases will be extended as required in future versions of the WRON-RM. Other artefacts will be required to manage the WRON and realise these key artefacts in practice.

Each information artefact will be created, managed and used according to the enabling design of the WRON infrastructure. The detailed specifications of each artefact will be developed by following a process that ensures both implementation feasibility and longer term effectiveness as the WRON evolves to meet more business needs. The interoperability framework described here outlines current best practice and future directions as well as a theoretical basis for achieving interoperability between the contributing sub-systems that the WRON will integrate.

5.2. Service Profile

The key artefacts in a Service Oriented Architecture are the definitions of the interfaces through which services interact. Service profiles encapsulate the ability of WRON stakeholders to publish and exploit data and information services, as well as processing service components that may be composed into information services.

Web Services Interoperability Organisation WS-I[24] states that general purpose technical specifications tend to be highly flexible and leave many options open to implementation choice. To be reliably useful, a set of pragmatic constraints, called a profile is agreed.

Service Profiles may be composed of a range of specific artefacts, for example Service Type, Data Product Specification, Query Model, Vocabulary bindings, Domain Model Elements,

[24] http://www.ws-i.org/

Implementation binding and Metadata bindings and should be as standardised as possible across the WRON.

Further elaboration of service profiles is left to the WRON-RM 0.1 validation phase.

Recent experience (e.g. AWDIP, OGC Gazetteer Profile (Fitzke 2006)) suggests that a useful data service profile will support data and metadata access in a consistent way: i.e. metadata may be accessed as a type of data object.

Each data service would support three distinct types of data object:

- Data (in rawest form available);
- Descriptive metadata (in most structured form available); and
- Statistical summary (domain set) for data instances, e.g. what time period or parameters are available for a water monitoring site.

Obviously statistical summaries are a type of metadata, but being dependent on object instances, this needs to be promoted to a particular information type. Where a predictable regime is implied, statistical summary may be implied and specified as a constant value in the service profile, possibly in the query model.

In future versions of the WRON-RM it is expected that the range of metadata artefacts included in service profiles will be augmented with additional metadata types, e.g. authentication, authorization, accounting and audit (AAAA) requirements.

The Service Profile mechanism provides an extensible container for a growing sophistication in the contract between supply, discovery and usage capabilities of the WRON.

5.2.1. Data Product Specification

The key artefacts in an Information Architecture are the definitions of the message content that is transferred between components and actors.

Version 0.1: To be developed for version 1.0.

5.3. Query Model

A query model is a description of the invocation query that can be used by a software client, or programmer creating a specific software tool, to invoke a specific instance of a data access service. It is a form of access metadata and is encapsulated by the service profile.

In this sense, the query model will bind a service type and data type to the actual content the user will interact with. Query models will typically bind a domain, such as a controlled vocabulary, or even the existence of data from another service, to an attribute of a data type sent as an input to the service.

NB: The terminology ‘query model’ may need to be revisited to avoid confusion with the meta-model of the implemented query languages (e.g. SQL or OGC filter).

Two examples of Query models are detailed below.

Example 1: (Dam Level Scenario)

Version 0.1: To be developed for version 1.0.

Example 2 (AWDIP)

AWDIP allows extraction of time-series data from an archive of observations. In general, this may be an arbitrarily complex and expensive query, so the AWDIP demonstrator phase defined a query model with specific bindings:

Query attribute	Data type property	Constraint
Sampling Station Id	Sampling Station Id	ID discoverable from Sites service
Observable Parameter	Parameter_id	ESCAWRI parameter dictionary
Start date	Sampling time	Expected date range
End date	Sampling time	Expected date range

5.4. Service Chain Composition

Service chains are created to produce information products (these may be considered virtual data services), or to produce visualisation outputs (which are just another form of data). Such service chains may be created by a user at run-time or by expert intermediaries, who then publish these as virtual data services.

Service chains may be created to implement existing service profiles, or be described automatically as part of a composition process as a new service profile.

Service chains may be visible to the End User if they need to control execution, but in general the intention is to simplify WRON interaction by encapsulating service chains as data services where possible.

5.5. Register (Meta-model)

Many implementation, governance or functionality evolutions will require specific information not specified within the enabling infrastructure. The infrastructure, however, must allow publication and usage of such information and thus needs to support a general register implementation. This is termed a meta-model because it specifies how a register may be specified, so that in turn it may be used to manage certain types of artefacts.

Focusing on the meta-model, the WRON enables new registers to be created as required. This is necessary because the requirements for ‘adequate description’ of services and their data will evolve over time as new functionality and integration requirements emerge. The principle of ‘adequate description’ is discussed in detail in Section 6.1.3.

Registers may contain:

- Specifications (types);
- Content (instances of these types); and
- Semantic value-add (relationships between registered objects and types).

The WRON-RM explicitly requires registers to be created to extend the information model of the WRON system. This ensures that WRON specifications and components are always accessible to stakeholders. In turn, this will enable future semantic exploitation to more rapidly build value added services.

Section 7.2.4 describes implementation options for registries, however it is worth noting that a flexible information model will be adopted for registers, with ebRIM being the basis for the key

initialisation phases of the WRON, where governance and management of artefacts is the prime consideration. This may be revisited as semantic descriptions become richer.

5.6. Stakeholder/Artefact visibility

This section provides a cross-reference between stakeholder visibility and the key artefacts. The cross-reference indicates the artefacts that are visible to each stakeholder, hence giving an insight into the degree of interest and involvement the stakeholder will have with that artefact. The cross-reference can be seen in Table 5-1.

Table 5-1 Stakeholder Visibility Cross-Reference

Artefacts	**Stakeholders**					
	1	2	3	4	5	6
Service Profiles	●	●	●	●	●	●
Service Chain Composition	●			●		
Query Model	●	●		●		
Register Meta-model			●			

Stakeholder List for ***Error! Reference source not found.***

1. End User Perspective
2. Data Provision Perspective
3. Enablement and Governance Perspective
4. Functionality Provision Perspective
5. Cross-Domain Business Integration Perspective
6. System Maintenance Perspective

6. Information Viewpoint - Interoperability Framework

The WRON-RM Information Interoperability Framework provides an approach for creating and growing the WRON in terms of more content and increased usability of data.

This framework concerns:

- The meta-model – how different information types interact to allow the WRON to work effectively; and
- The methodology for creating information elements so that the WRON continues to work effectively.

6.1. WRON Meta-model

Meta-models provide the framework for the design of the domain- and metadata-models, and thus allow multiple information packages to be created using a common pattern, and hence extend, include or relate to each other. The WRON meta-model framework distinguishes between:

- Common patterns and components (abstract models that encapsulate common behaviours, and concrete components that are used in multiple domains);
- Domain Model (conceptual model of content for a business domain);
- Metadata Model (what descriptions are necessary); and
- Implementation Model (how to build with particular technologies).

Support for common implementation approaches at each of these levels allows for WRON stakeholders to easily understand how the WRON works and how to build components that will interoperate, without having to be familiar with the entire WRON and each component.

The approach described here ensures that all information artefacts implemented in the WRON can be used according to the business requirements identified in Section 4.

Critically, this includes integration between related infrastructures (for example, the WRON and state-based systems or the WRON and related business domains) as well as cost-effective implementation by commonly available technologies.

The ISO 191XX standards (summarized in Appendix A and B) provide the framework for the domain and implementation models, and parts of the metadata model for the WRON-RM v1.0

6.1.1. Common patterns and components

If common patterns and components are used in the creation of information packages these may be designed to re-use or profile each other in a consistent manner, and software implementations may be shared across business domains.

In the framework defined by the ISO 191XX standards the following are provided (see Appendix A and B for references):

1. Standard patterns for the primary geospatial information types – features and coverages (an upper ontology);
2. A standard treatment of coordinates and coordinate reference systems;
3. A harmonised treatment of geometry, topology and time; and
4. A standard model for registration and profiling.

In addition, the Observations and Measurement pattern has been designed consistent with this framework. It provides standard access points to metadata related to a measurement, regardless of business domain and hence of observed phenomenon (see Figure B-1). The associated Sampling Features provide an access method for observation collections which is common to multiple scientific domains (see Figure B-2).

Other common patterns are published in the WRON as modular model packages, in accordance with the modelling methodology described here.

Version 0.1: An enabling suite of common patterns will be identified, tested and published for Version 1.0.

6.1.2. Domain model

A *domain model* describes the concepts and relationships that support discourse within a domain, i.e. the language of the community. This should be expressed using the terminology of the domain (e.g. dam, catchment, stream) in preference to a geometric or data-capture abstraction (e.g. point, polygon, scene).

A domain-model or ontology underlies any meaningful applications within a domain. Hence, a model for data-resources is itself a key resource required to define the system. For practical use of a model in a system it must be accessible (published) and stable (governed). [Data resource semantic representation]

Instance data transferred within the architecture must then conform to the model. This implies that serializable implementations of the domain model are available. [Data resource syntactic representation]

The domain model provides authoritative definitions of concepts within the domain. When all services are consistently bound to the domain model, semantic interoperability is achieved. Profiles, as containers of this binding, must use concepts that are found within the domain model. Parts of the domain model may be under governance by a responsible body (such as a standards body) and it makes sense that the true form be kept there at the ‘point of truth’. The domain model will need to reference these parts and there is a versioning requirement as the ‘point of truth’ may change in a way that invalidates the domain model and requires manual intervention to fix.

The ISO 191XX standards provide the basis for formalising the domain model for the WRON-RM v1.0. The representation of the domain model will be a taxonomy of concepts, encoded in UML according to the rules in ISO 19103, ISO 19109 and ISO 19110. In a future version of the WRON-RM, the automation and extensibility of the domain model may be improved by incorporating into the methodology more expressive representations of the domain model, such as ontology.

6.1.3. Metadata model

Pragmatism

Metadata (‘adequate description’) is an evolving and complex requirement that must be aligned with both external standards and business processes of WRON stakeholders. It is not necessary to finalise the decomposition of metadata into standard forms to realise the WRON, as long as the ability to add new types and revisit implementation of existing types is built into the WRON-RM. Where the WRON will rely on certain types of metadata, an implementation binding will be provided, along with a roadmap for standardisation.

Metadata requirements

The WRON requires adequate description to realise the discovery and usage business requirements identified in Section 4. This information can be broken down into a suite of

common metadata components that can then be implemented using a small set of standards, and reused across multiple applications.

The WRON will require an evolving set of metadata as functional capabilities are added to the network. Thus, the WRON-RM will define a framework for managing metadata.

This framework will mandate:

- Clear governance and lifecycle management of all metadata;
- Metadata objects defined in WRON Use Cases (in particular those identified as part of a contract between stakeholders) will have a normative form, specified by a WRON profile (defining conceptual and implementation bindings), and over time, be aligned with relevant external standards;
- References to other objects will conform to a small set of WRON-endorsed implementation approaches to ensure WRON clients can expect to traverse such references as required;
- Use of common vocabularies (by reference) rather than duplication;
- Replication of content will be avoided, and will retain reference to point of truth where it is deemed desirable to insert information managed elsewhere into an artefact;
- Additional metadata artefact design will conform to external standards where possible, to allow interoperation and integration of the WRON and other infrastructures;
- Metadata may be applied to an external resource by reference, so that the WRON can dictate how a particular resource can be used within WRON compliance expectations;
- Where available, metadata in the WRON registry will override metadata supplied directly by a resource. As user profiles are developed in later versions this may be extended to allow user overriding of metadata; and
- Where possible, metadata in the WRON registry will be harvested directly from the authoritative source, and augmented according to WRON compliance requirements.

Metadata purposes

The Orchestra RM2.0[25] provides nine purposes for metadata and the WRON-RM adds a tenth, namely Visualisation. These purposes are a useful classification of the metadata in the WRON. The domain model will conceptually define metadata associated with these purposes and the service profile will bind them to a particular implementation. As such the metadata definitions will be brought together in the service profile.

As with all things in the WRON, the actual metadata instance should reside at the ‘point of truth’. The service profile will therefore contain bindings describing how the metadata can be retrieved from the services. Since metadata can be retrieved essentially like data this may even include a query model for each type of metadata.

Version 0.1: Not all the purposes will be covered in the WRON 1.0. After the implementation, testing and validation the following will be specified:

Description of Metadata Purposes

Discovery

Discovery is aided by the entire service profile and domain model. During discovery clients will be able to pose queries across service profiles and domain models because they are stored in registers, and so are accessible to search. One type of metadata in particular is used primarily for discovery – summary data.

[25] Reference Model for the Orchestra Architecture (RM-OA), 2007-01-31, section 8.4.2

Summary data is best described as an operation over the content of a data service. This operation could be a simple operation on the service serving the content, a registry service that stores the summary data or even a service change composition.

As such, the summary data operation will be represented by the Access metadata described below. This is true of most metadata, except of course the Access metadata itself.

Access, storage and service invocation

This is the most critical part of the metadata that makes up the service profile. Following the principle of 'adequate description' it must say enough about the service so that it can be invoked in a reliable fashion while not overly burdening the service provider with description. At least five levels of description can be identified:

1. A name;
2. A name, service type (e.g. WFS);
3. A name, service type and description of the inputs and outputs (IO);
4. A name, service type, IO and preconditions and effects (PE); and
5. A name, service type, IO, PE and a workflow to accomplish the multistage process.

Level 1 would require every service of the same name to be identical. This would be insufficient and developers would try and encode information in the name to overcome the lack of descriptiveness. A more appropriate level of description is Level 3. At this level, the inputs are described in terms of a query model, which in turn uses terms from the domain model, and the bindings to a particular implementation (such as WFS or SOS). The output could be of fixed type; however in the case of data services the query determines the output type so it is best to simply use the data schema as the output. Outputs require bindings to terms from the domain model.

Integration

Integration is available on two fronts. Firstly, because profiles are predefined, any implementations of the profiles will have known semantics and syntax. The services can therefore be integrated. Secondly, because all services are related to the domain model an advanced semantic interpreter should be able to automatically integrate services. This is a topic for future research.

Interpretation

By relating everything to terms of the domain model concepts have agreed meaning and can be used consistently across the domain. There are two contracts, beyond the infrastructure still required to ensure the WRON works. Firstly, the providers of data and services must examine the definitions of domain model terms and ensure their services properly reflect those definitions. Secondly the users must examine the definition of domain model terms and ensure they are using the information in the right context.

User profiling

Users, and enterprises, will have their own preferences and terminology which may be different from the terms in the domain model. User profiling will cover mechanisms for ensuring user requests are translated into the community's terms in the domain model and that results are rephrased in terms the user understands.

Authentication, Authorisation, Accounting and Audit

Authentication, Authorisation, Accounting and Audit (AAAA) are out of scope for the first version of the reference model. Audit enables transparency and traceability. It will be a key feature of the WRON. Some traceability is already in place via the Observation and Measurement (O&M) pattern. The O&M information model encapsulates well the provenance information associated

with data (along with a detailed description of the procedure). The use of O&M and audit generally will be expanded in future versions.

Quality Control

Version 0.1: To be completed in a future version.

Transactions

Version 0.1: To be completed in a future version.

System configuration and management

Version 0.1: To be completed in a future version.

Visualisation

Some data providers helpfully provide visualisation advice along with their data. This is often in the form of standard symbology to use. In many cases, communities will mandate particular presentation formats for information. The WRON must support multiple alternative representations for data products. This is closely related to user (or community) profiles.

The WRON must also support the adoption of a WRON-registered presentation option by external stakeholders. This may require transition of governance arrangements and lifecycle evolution (move to a different stage). Transparency of governance is critical to meeting this function.

6.1.4. Implementation Model

Standards

The WRON will be built on top of suitable existing standards (see Section 1.2.1). These standards include generic cross-domain components, and compatible existing subject-specific standards that satisfy particular WRON requirements. The standards to be adopted are those where evidence of implementation and capability within the sector can be identified.

Many WRON artefacts will be implemented without standards to guide all aspects of the implementation, and this is expected to result in feedback to appropriate standards bodies.

Version 0.1: The WRON RM 1.0 will be based on:

- *ISO TC211 Geographic Information (UML, Feature-model, Registration model, GML)*
    - *implies OGC implementation specifications (WFS, SWE);*
- *W3C Web Services (WSDL-2, OWL);*
- *OASIS (ebXML Registry implementation); and*
- *Web Services Interoperability Organisation (WS-I).*

It should be noted that there is formal liaison between these bodies already, which the WRON will exploit to maximise consistency between different aspects of the Reference Model.

Representations

Information-viewpoint resources will be expressed as serialised artefacts which may include, but are not restricted to, the following:

- OWL documents for ontologies;
- XMI documents for UML models;
- XML Schema documents for data types;
- DDL expressions for table schemas;

- SKOS documents for dictionaries of terms;
- XML documents for data instances;
- XSLT documents for portrayal rules;
- PDF documents for interface specifications;
- WSDL documents for interface types; and
- OGC Capabilities documents for service instances.

These artefacts provide a description of system-level components, complementing artefacts that represent data instances. This 'metadata' provides configuration information to support server/client binding and service composition.

6.2. Modelling methodology

The WRON-RM is being developed primarily as an integrated UML model, focusing mainly on use-cases, packages and classes. The model template being used is known as HollowWorld, and was previously developed to provide an observation-oriented, ISO conformant environment. HollowWorld includes the complete ISO 191XX Harmonized Model, and a representation of key OGC Sensor Web Enablement capabilities.

In future versions the distinction between conceptual vs. implementation models will be developed.

The following section lists and describes aspects of the modelling methodology that will be used as part of the WRON-RM v1.0.

6.2.1. Layered models

Models will have multiple meta-levels and ways of moving between these levels are required. Two possible ways of transitioning between meta-levels are:

1. Meta-level transitions: Both the domain model and the implementations may have inheritance. This allows transitions from the abstract to the concrete.

Version 0.1: This section will be completed in version 1.0.

2. Specialisation: The leaves of an inheritance graph are more specialised resulting in transitions from general to specialised.

Version 0.1: This section will be completed in version 1.0.

6.2.2. Service Profiles

While the domain model defines the semantics of a business domain, the service profile declares how the data, metadata and processes in the domain model will be implemented in the WRON. Definition of service profiles, and compliance with them, provides services with well known semantics and syntax that can be readily integrated.

6.2.3. Governance

A strong governance model is necessary to maintain order and usefulness in a system like 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.

Furthermore, the semantic value of each definition used within the WRON will ultimately be understood in the context of its governance: who defined it and why. Thus, indicators defined to meet a legal requirement and trends extracted from raw data may be identical, but the indicator will be discovered and used, and carry metadata relating to its use.

Registration

Orderly registration is required by the WRON service model to ensure that a service consumer can access and bind to services that are required. It is also critical that the registration register is appropriately managed to ensure currency of information.

There are several roles of registration:

- Publishing: Make information available so it can be used;
- Discovery: Allow the location of information and services to be found; and
- Governance: Provide a control mechanism to ensure quality and consistency.

Profile Implementation

Profiles are combinations of concepts and implementation bindings which represent a useful functional unit to a community. The publication of a profile, and subsequent conformance to this profile, will be a key governance mechanism to ensure End User expectations can be met.

Service Level Agreements

Service Level Agreements (SLAs) will determine the basic requirements for data availability so that the WRON can meet its business requirements. This may be a single, or suite, of options, but must be readily understood by End Users and Service Providers alike.

This means that SLAs are key artefacts that form part of the WRON implementation, and not a register that records many variant implementations.

6.2.4. Identity

The WRON aims to support effective access to the broad range of information required to model, monitor, understand and manage water resources. Water resources *features*, such as a river segment, supply network, storage or catchment will have many different characterisations and related information products. Identity (of feature instances) is thus an important information concept in the WRON. In order for information to be usefully integrated the identity of entities referenced by different custodians must be reconciled.

Approaches

Identity management and resolution of pre-existing alternative identifiers is a complex problem with no definitive solution. Several approaches exist however, which are not mutually exclusive, and may be applied to different sub-sets of the WRON information domain:

Master register

One approach that may be used is to delegate identity management to a central organization, and each data provider undertakes to use the assigned identifiers. This is the approach used by the UK *Digital National Framework* where object identifiers (TOIDS) assigned by the Ordnance Survey are used by each participating agency.

This approach can be applied to specific sub-domains. Candidates for such an approach include:

- Hydrologic features;

- Supply networks;
- Storage; and
- Catchments.

It must be noted that often a real-world feature plays multiple roles, with a different characterisation for each purpose. For example a river may have a historical course, used for cadastral purposes, and current physical properties, used for hydrological modelling.

An important function of the WRON is to provide a framework that allows such complex problems to be efficiently resolved, and then effectively exploited. The 'common model' approach described below may be required to resolve identification design.

Persistent local unique identifiers

Where an organisation delivers data, the identifiers must be stable – the identity and the format of the identifier must remain stable, to allow third party systems to create references that remain valid. Accordingly, the use of identifiers by external partners is a critical Use Case that must be part of each service (data product) specification. Metadata for the identification scheme in use must be published in a simple, standard format.

In this methodology, mappings between identifiers can be introduced to the system on an as-needed basis. This requires development of a common id-mapping artefact design, and establishment and governance of suitable a register of mappings, since each mapping needs to be discoverable, and represents a significant investment that must not be replicated.

Common model

When different aspects of the same real-world feature are represented in different systems, it is unlikely for the identity to be consistently managed.

For example, geometry-centric implementations (such as ArcHydro) tend to use the geometric representation as the essence of the existence of the feature, the 'identity-carrier'. However, geometry changes with changes in geomorphology, with better surveying, and with generalisation for processing/transfer efficiency.

Other systems may use different named attributes for identity, often with identity being a qualified (composed) set of attributes.

The common model approach provides for a separation of concerns into a common conceptual model and multiple implementation models. The common conceptual model carries identity (i.e. the custodian of the conceptual model determines the identification process for the feature). Each implementation model will be responsible for mapping its implementation of identity onto the common model.

Semantic Interoperability can only be defined at the conceptual model, so the resolution of identity management issues becomes an integral part of service design, with the WRON supporting and propagating the resolution.

This methodology provides for a single, clear and simple governance path for identity management, and a responsibility for this as part of product design.

Purpose Orientated Market Driven Common Model

Similar to the common model this form adds two flexibilities:

- The choice of identity mapping can be determined by the purpose for which that mapping is to be used. For example at different spatial scales there may be different identity mappings; and
- Market forces are used to provide different mappings of varying quality which may be automatically or manually obtained (typically for a completeness/accuracy trade-off).

There is a requirement for a strong governance regime for the mappings.

There is some experience in the use of each of these approaches, and others may be possible. It is unclear at this stage whether all approaches must be implemented simultaneously, or whether a single mechanism can be safely prioritised.

Semantic interoperability between value-adding processing and data supply tends to suggest that common model approaches underpin scalability of any master registry or cross-reference management approaches.

Validation of a coherent approach to identifier management will be an important part of the preparation of the WRON-RM v1.0. Validation requires significant cross-domain data model implementation, with realistic quantities of information to assess the inherent challenges.

6.2.5. Baseline Information Models

The WRON will require implementation of certain types of information artefacts to function. This set of information artefacts is termed the Baseline Information Model because it provides a targeted set of components that must be designed and implemented in deployed systems for the WRON to exist. Each business application will extend this baseline, and hence the WRON will be the aggregation of a growing set of applications. The existence of a coherent baseline means that the utility of the WRON can grow beyond the sum of the individual parts.

This section outlines the conceptual framework provided by the adoption of the initial set of standards suites and describes both the building blocks and the best practice methodologies for using them.

Web Services

Web Services Interoperability

Service Profile implementation will provide WS-I conformant compliance profile artefacts.

ISO 19100 framework

Spatio-temporal context is critical to most expected WRON Applications, so the WRON-RM must pay particular attention to geographic information (GI) aspects of the information elements. A relatively mature methodology for specifying both information and services within GI has been developed by ISO Technical Committee 211 – Geographic Information (ISO/TC 211), partly in collaboration with the Open Geospatial Consortium (OGC). The ISO work is expressed as a set of International Standards and Technical Specifications in the ISO 19100 series. A summary of the ISO models and standards is provided as Appendix A and B to this document.

The WRON has several powerful drivers to adopt this framework, where it applies:

1. This is an accepted set of best practices that represents a significant amount of effort;
2. The ISO framework is explicitly implemented by OGC specifications, and hence by most significant technology providers (to some degree, but at least in principle);
3. The framework is ‘implementation neutral’ and avoids any problematic technology lock-in, and there is unlikely to be any alternative governance structure for an implementation-neutral solution;
4. The ISO suite exists as a series of model packages that can be directly imported and used;
5. Widespread adoption by ANZLIC, INSPIRE/ORCHESTRA, UNSDI/GSDI umbrella programs; and
6. International relevance of the skills and solutions developed for the WRON, which is important for most participating enabling partners.

Cross-domain standards

The ISO 19100 standards 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 GI applications. Applications that share these components are interoperable at this level.

Models have also been developed for more specialised applications, such as Natural Environment, and Sensor Web. Where these have also been developed using the ISO domain modelling framework they form an interoperable standard set.

OGC Sensor Web Enablement

The OGC Sensor Web Enablement (SWE) initiative has defined a set of service interfaces to provide a generic access method to sensors and observations. The SWE service interfaces include Sensor Observation Service (SOS), Sensor Alert Service (SAS), Sensor Planning Service (SPS) which follow the OGC Web Service interface patterns established by Web Map Service (WMS), Web Feature Service (WFS) and Web Coverage Service (WCS).

Observations and Sampling

Underlying the SWE interface standards are models for Observations and Measurements (O&M), Sampling Features, and Sensors (SensorML). These have been shown to be effective in describing observational data across a variety of domains concerned with the natural environment, and have been used in the Australian water resources context in the AWDIP demonstrator.

Water-domain standards

There are a number of water domain initiatives that are of interest to the WRON. Two which are of particular interest are the US CUAHSI/HIS and the Australian AWDIP project.

CUAHSI/HIS/WaterML

Version 0.1: To be completed in version 1.0.

Note that WaterML has recently been submitted to OGC as a potential Discussion Paper.

AWDIP

Version 0.1: To be completed in version 1.0.

Note that AWDIP provides an O&M testbed.

Canadian water project

Version 0.1: To be completed in version 1.0.

6.3. Worked Example

6.3.1. Dam level Scenario

This section uses the modelling methodology defined in Section 6.2 to explore and then list the information artefacts and flows that are necessary to support the Dam Level Scenario. The artefacts are identified by the solid boxes and the flows by the dashed lines in Figure 6-1.

Version 0.1: This scenario is not currently fully developed and will be competed in Version 1.0 of the WRON-RM.

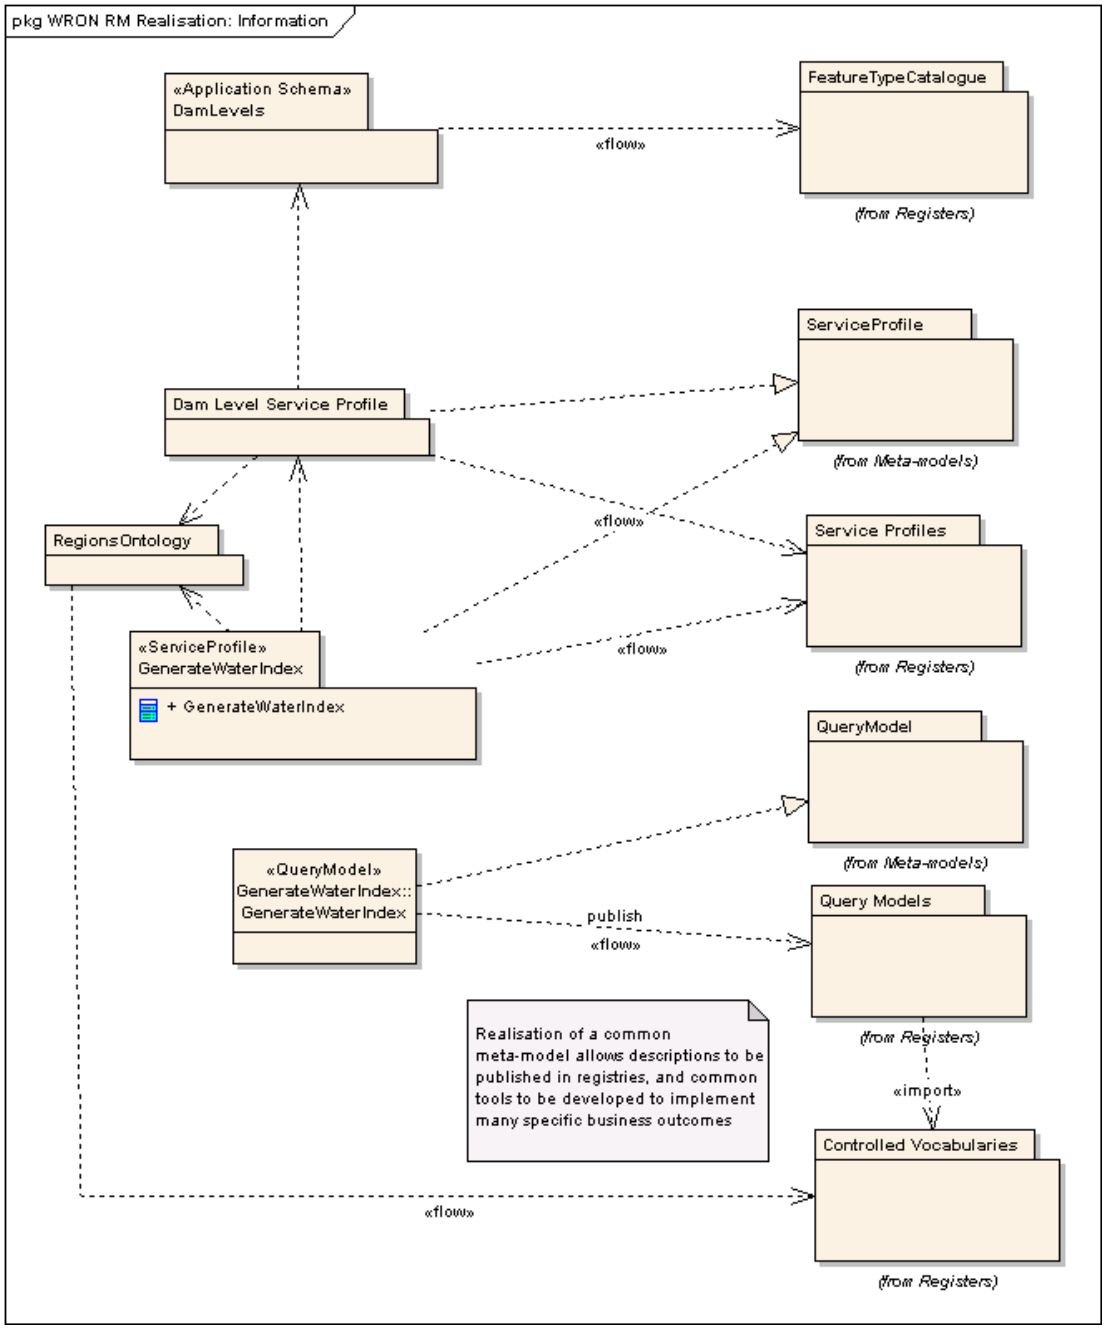

Figure 6-1: Dam Level Scenario Information Artefacts

Dam Level Scenario Information Artefacts:

- Dam Levels Application Schema;
- Dam Level Service Profile;
- Generate Water Index Service Profile;
- Regions Ontology;
- Generate Water Index Query Model;
- Feature Type Catalogue;
- Controlled Vocabularies;
- Registers Query Model;

- Meta-Model Query Model;
- Registers Service Profiles; and
- Meta Model service profile.

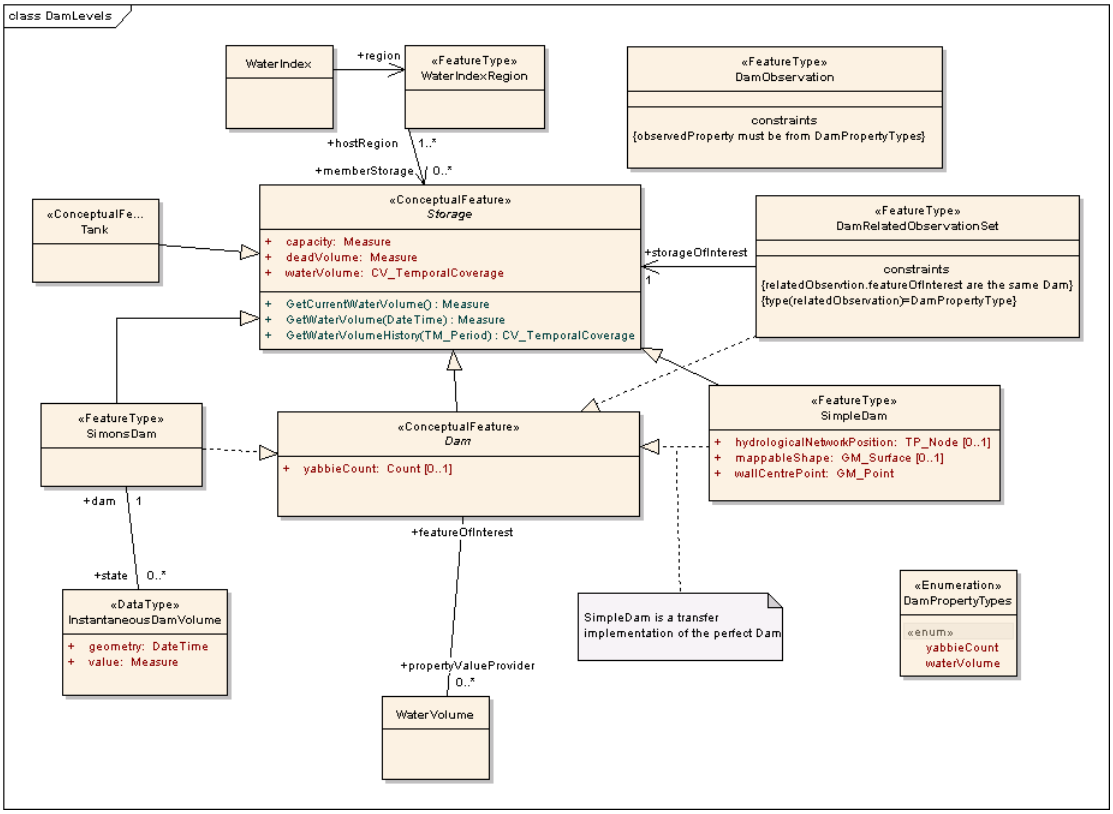

Figure 6-2: Dam Level Conceptual Schema

Information Model for the Dam Levels Scenario

Version 0.1: To be completed in Version 1.0

7. Computational Viewpoint

The computation viewpoint defines the components in the system and how they interact with each other to form a computable network.

7.1. Overview

The Computational Viewpoint of the WRON-RM describes the functional breakdown of the system into components and the interfaces between them. This means identifying the set of core services that comprise the infrastructure and the types of services that will be deployed to realise specific business functions.

The overall goal for the WRON is to connect client software with services to deliver the information products they need in a way they can understand. This information will come from potentially disparate sources who publish their services into the WRON using different standards.

This means that the computational framework is dictated by the business requirement for multiple stakeholders in both the data delivery and data exploitation aspects.

7.1.1. Extending Best Practice

Enterprise service integration is not new, however the WRON offers three things that set it apart from other initiatives:

1. **Client closeness:** Enterprise integration techniques often expose either their own interfaces or that of the services they integrate. The WRON aims to deliver information to the client in a form, and through services, they can deal with;
2. **Multi-standard:** Enterprise integration systems generally deal with one standard, typically their own. The WRON integrates a limited, and changeable, number of different standards which have independent governance; and
3. **Semantic integration:** In addition to simple syntactic integration, the WRON-RM v1.0 will define the semantic building blocks for future scalability and improved user support through semantic ‘mediation’.

7.1.2. Specification methodology

The WRON will support publication of service type specifications as a suite of service profiles. Detailed specifications for each component type will thus be provided in a way which can be further profiled into a directly implementable profile for a particular business need, e.g. AWDIP.

The WRON-RM does not need to specify in detail each component type, but the expected functionality is identified from the business requirements perspective, as well as key enabling functions that will require implementation as infrastructure**.** The WRON-RM v1.0 will thus include the published service profiles for critical components.

7.1.3. Point-of-Truth vs. Data Warehouse pattern

A traditional approach to data integration is to create a data warehouse by importing the information from sources, hopefully on a regular basis (e.g. monthly). During the importation phase the contents are reconciled to the local schema. This has many problems including:

1. **Currency:** The warehouse is expensive to update regularly, particularly if it contains a variety of data from a variety of sources;
2. **Transparency:** Data reconciliation may introduce interpretations or errors, and this is problematic to load onto a single agency; and
3. **Control:** The creating agency will lose control of, and be unable to monitor the usage of, data by sending it on.

The WRON aims to ensure data is made reliable (useful) at the 'Point of Truth' and request only what is needed and when it is needed. This ensures the data is always up to date and the creating agency (custodian) maintains control of that data. The agency can withdraw access at any time or even apply corrections to the data. If performance is an issue, then limited caching, with the agency's approval, can be performed by the middleware.

If required, a data warehouse can be added to a Service Oriented Architecture (SOA) at any point, to support specific service level agreements. With a data-warehouse centric solution, new data and services cannot be reliably added without modifying the warehouse.

7.2. Key Artefacts (Components)

This section identifies the key services and components that are required to realise WRON business functions. These components have been identified by analysis of the system Use Cases. They are the services that users are likely to have direct interaction with. This version of the WRON-RM limits the scope to identification of services and components, in the WRON-RM 1.0, the services and components and their interfaces will be appropriately defined.

7.2.1. Data Access services

Data access services are services from which data originates, as opposed to processing services which output a modification of the input data. The difference between these types of services is somewhat blurred, with data services often doing some processing and processing services using some local data sources.

The most distinguishing feature of data services is that they are described most often in terms of a query model and data schema rather than inputs and outputs. This is because the output is a function of the input query model and the data schema. Simple services (described in Section 7.2.3) are another type of data service.

The WRON will support those implementations of data services which have registered profiles. Profiles are described in the information model in Section 5.2.

The WRON-RM v1.0 will define the AWDIP profile and hence the WRON will support OGC's WFS implemented in accordance with the profile. It is expected that other yet to be determined profiles will also be defined in WRON-RM v1.0 and that this list will expand over time as more profiles are created and verified.

Metadata Services

Metadata is data and therefore can be accessed with a data service. All the metadata listed in the information model could, if necessary have services to provide access to them and hence makes no distinction between metadata and data services in the Computational Viewpoint. For example summary data (metadata) can be provided by a processing service which applies an operation to a data set or a registry service which extracts the information from a capabilities document.

7.2.2. Processing Services

Processing services will be required for a broad range of purposes, from complex modelling to simple type transformations to report generation. They are distinguished from data services in that they are normally defined in terms of well constrained inputs and outputs. See data services (Section 7.2.1).

Like all services, the WRON will support those implementations of processing services with registered profiles. e-Water catchment toolkit (TIME) models will have a representation as WRON processing services but their profiles are yet to be defined.

7.2.3. Simple Access

Simple access services have very simple interface descriptions and therefore have the advantages of simple definition and consumption by a client. These services generally take a few simple concepts as input and return a fixed output type. For example, a service could be deployed over rainfall data that accepts a town ID and a start and end date and returns a time series of rainfall. Simple Access services are particularly useful in the WRON in the application services area where they enable applications like Microsoft Excel to interact with the WRON.

7.2.4. Registry

Registry services provide access to the WRON registers. Governance of a registry service, to ensure register data currency is a core concept and requirement in the WRON. Custodians of services will, often, keep their own registry and this represents the point of truth for information pertaining to those services. Within the WRON there will be central registries which will refer to the custodian registries as well as maintaining registers common to the WRON as a whole. The synchronisation of information between these registries will be a key technical issue to be resolved in the WRON-RM v1.0.

Registry services will have a profile associated with them and the WRON will support those implementations defined in profiles.

The initial implementation will support:

- ebXML registry information model (ebRIM) for flexible object and relationship management (implementing ISO 11179 Metadata Registry meta-model); and
- Version control (e.g. Subversion) allowing common technologies to access and control WRON specifications.

Common repositories may be accessed by multiple protocols, and content may be mirrored between multiple repositories to provide alternative protocols and query capabilities. It is expected that mirroring between commonly available technologies will be trialled. This does not prejudice consolidation of functionality into a single registry component later.

Other registry implementations (such as UDDI and EPCIS Discovery service) will be explored in future versions.

7.2.5. Semantic Registry

Reference is made to a semantic registry now so that readers can be made aware of the future directions of the WRON-RM. Registries are searched by posing queries in much the same way as querying a database. A semantic registry will take semantic queries in a knowledge language such as description logic. This will enable a more goal focused query which can reason over the domain model and locate an appropriate service. The semantic registries might not only find

services better, but also generate queries for them or compose service chains and automatically publish these as useful new services.

Advanced clients, in future, will interact directly with the semantic registry. A more likely case in the medium-term is clients binding to a standard registry proxy which translates their ordinary registry request into a semantic request based on the user profile.

7.2.6. Supporting Service Types

In addition to the components identified above, being those services (components) that stakeholders directly interact with, the WRON will require additional service components that assist in the execution of functions. This section defines the supporting services that will be required for the WRON.

Wrapper (Adaptor) Services

Where a stakeholder wishes to publish a service that does not comply with any profile they have three options:

1. If their service conforms to a profile's conceptual model they need only supply an implementation binding for their service and have the profile governance body approve it;
2. Apply to the profile governance body to have their service accepted as a new profile; or
3. Write a wrapper (adaptor) over their service to make it conform to an accepted profile.

Data Warehouse

A data warehouse is a useful construct where rapid response (and complex database operations) are required. Traditional data warehouses have the problems indicated in Section 7.1.3, however where a data warehouse augments a 'point of truth' system like the WRON some of these are mitigated. Updates can happen at any time because the sources are online and the profiles provide known, stable aggregation points. A data warehouse integrated with the WRON should be seen more like a distributed database with the warehouse being a data heavy node.

Mediation Services

Mediation services look like any other service and comply with profiles. They are special however because they aggregate or otherwise value-add one or more services. Without mediation services clients would only be able to access services published by the 'point of truth' service providers and would have to manage working across multiple services. Instead mediation services perform common compositions for clients.

There are several ways the WRON can assist the creation of mediation services. These are covered in the Computational Framework section below.

7.3. Computational Framework

The WRON is a Service Oriented Architecture (SOA), without being bound to the current W3C SOA implementation stack. It assumes that functions may be decomposed into separate services that may be deployed by different organisations within a network environment.

The WRON assumes a growing number of processing and mediation services will be deployed against a relatively stable, incrementally expanding, suite of data services. The WRON-RM v1.0 will require, at the very least, the ability to deploy new instances of service components, and new types of service components.

The WRON Computational Framework provides a methodology for maximising commonality of approaches so that the cost and barriers for exploiting new functions can be minimised.

7.3.1. Role of standards

Business domains develop standards to allow them to have a degree of interoperability amongst members. Generally these standards achieve a degree of syntactic interoperability. They then rely on sophisticated clients and users to understand how to combine the information from multiple sources. This is workable in very tight communities but the WRON is dealing with hundreds of different agencies, all familiar with doing things their way. Not only do the agencies have different understanding of concepts, but they also often use technologies from different standards.

7.3.2. Service Framework

Profiles

The key enabler in the service framework is the publishing of profiles in conjunction with a common domain model. Profiles, as described in the Information Viewpoint, contain the conceptual model of the service and its binding to implementations based on supported standards. The concepts used to describe the service are taken from the domain model. With all supported services being defined in concepts from the domain model they share common semantics. This will in turn allow services to be exploited in a semantically consistent way.

Version 0.1: Examples of this will be developed for the WRON-RM 1.0.

A corollary of the separation between conceptual and implementation services is that facets of a single service concept may be provided by separate physical services. For example a WFS service may provide data but its summary data may be provided by a query over a registry service.

Client Interaction

In principle clients should interact with the service with the interface they want, the schema they desire, and using terminology they understand. In practice interaction is limited to those services, schema and terminology supported by profiles published in the WRON. The onus is on the WRON community to develop useful profiles and implementations. The profile system is not tied to any particular standard so clients must lobby to have their standard supported.

Future work in semantic mediation (see Section 7.3.3) may improve the flexibility.

7.3.3. Mediation

Mediation services are a mechanism for providing an interface the client understands and deal with the detailed interactions with source services. Mediation service providers may take advantage of facilities provided by a Registry, Semantic Registry, other services and provider knowledge. Below are potential ways mediation could be provided. Each of these will need to be tested in an implementation phase to identify the most useful forms of mediation.

Types of mediation:

1. Service provides mediation, based on code embedded semantics, and integrates *n* services by looking for them by profile. The service deals with the implementation interface of each source service;
2. The WRON provides invocation facilities, for supported standards, which the service uses explicitly to convert *n* services, of the same profile but different implementation standards, into a common format. The service integrates them and then uses a

translation service to convert to the client format, in accordance with the mediation services' profile;

3. The WRON provides semantic composition facilities which allow a service to request a response from data sources matching a profile, with semantic and format consistency;
4. The WRON provides composition facilities which allow a service to formulate a request in terms of the desired result and the composition facility will construct a workflow;
5. The service presents itself as a processing service matching a profile. An expert user uses the composition environment to construct a new profile which uses the value-added service as part of the workflow; and
6. An expert user presents a semantic goal which includes what the new profile should look like and the WRON constructs a service and profile to meet that goal which in turn invokes a workflow of services.

Semantic mediation is an extension of syntactic mediation. Syntactic mediation converts data from one format into another. Semantic mediation adds extra processing steps so that the meaning of the concepts is consistent between domains, even if the format is different.

7.3.4. Semantic framework

Profiles bind the conceptual model of a service to implementations. The conceptual model is defined in terms of the WRON community's domain model. If this domain model is sufficiently rich and the services adequately described, machine reasoning tools may be used to aid discovery, mediation and service chain composition. These functions can be housed within a semantic registry and accessed via a logic-based query. For service chain composition across implementations from different standards the semantic registry may also need to provide invocation assistance.

Development of a semantic registry will enable rapid development of service profiles and mediation services. This is an active research topic for future WRON versions.

7.3.5. Dissemination and persistence of metadata

System metadata needs to be accessible at the point of use, but may in fact be published and collated in many places in the architecture. A significant architectural concern is where and when exchange of metadata occurs.

In closed systems this is often performed at run-time, involving some form of dynamic configuration of the client. For example, GIS data assimilation usually requires a human operator to inspect the incoming data, or possibly a persistence schema, and then configure a mapping to their local data model, or select fields to use (store, display). This places no specific expectations on the service provider regarding the information model for the data product that it provides. It is merely necessary that the service provide a description of its interface and data schema, which are instance-specific and therefore potentially unique. A unique information model is likely to be at least partially incommensurate with the client's requirements. It is likely that there will be insufficient clarity of definition to safely exploit data without extensive prior knowledge of the data. This strategy thus reduces the probability of consistent behaviour and analysis within an application domain, and is clearly not scalable. It is also an effective barrier for data being turned into timely information for End Users concerned with the management of water resources.

The principal alternative is to publish the system description in advance, and for participating components to subscribe to these definitions. This allows clients to be pre-configured with rich functionality corresponding to the richness of the specific information model and service interfaces. Conforming (data) services will generally require mapping to the public model from a generally private (data) source, though this function may be positioned in various places (e.g.

database view; within a service component tightly bound to the database; in a transformation wrapper; by a separate mediator service). The key requirement, and challenge, is a governance- and publication-process for the service metadata, and an agreement by participants to conform.

8. Engineering Viewpoint

Version 0.1: This section will not be delivered as part of WRON-RM 0.1 as it relies on completion of the Information and Computational Viewpoints. It will be complete for WRON-RM 1.0 and will cover such areas as: versioning, quality of service, exception handling, authorisation, authentication accounting and auditing.

9. Technology Viewpoint

Version 0.1: This section will not be delivered as part of WRON-RM 0.1 as it relies on completion of the Information, Computational and Engineering Viewpoints. It will be complete for WRON-RM 1.0 and will cover such areas as: specific technological solutions, existing technologies, technology engagement strategies and implementation gaps.

A. ISO 19100 Framework

A.1. Overview

The basic model is described in ISO 19101:2002 *Geographic information—Reference Model*. The following three figures extracted from ISO 19101 summarize some of the key concerns.

Figure A-1: From reality to conceptual schema (ISO 19101 Figure 4) indicates the context for information modelling, in the form of a 'Universe of Discourse' in the real world (i.e. an application domain, such as Water Resources Observations) which is summarised in a Conceptual Model for the domain (comprising a set of concepts and relationships used by practitioners within the domain). A formalism (notation) may be used as the basis of a Conceptual Schema Language, through which the Conceptual Model is represented as a Conceptual Schema.

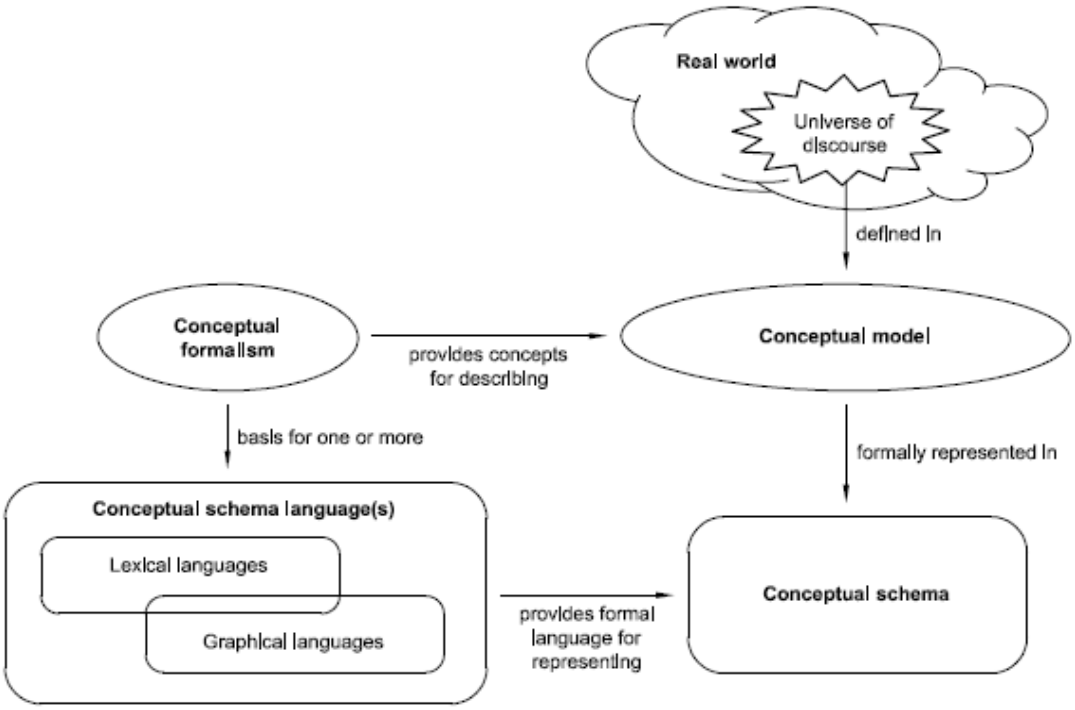

Figure A-1: From reality to conceptual schema (ISO 19101 Figure 4)

Figure summarizes a scope for data models within an Application schema, in terms of a meta-model based on positioned features.

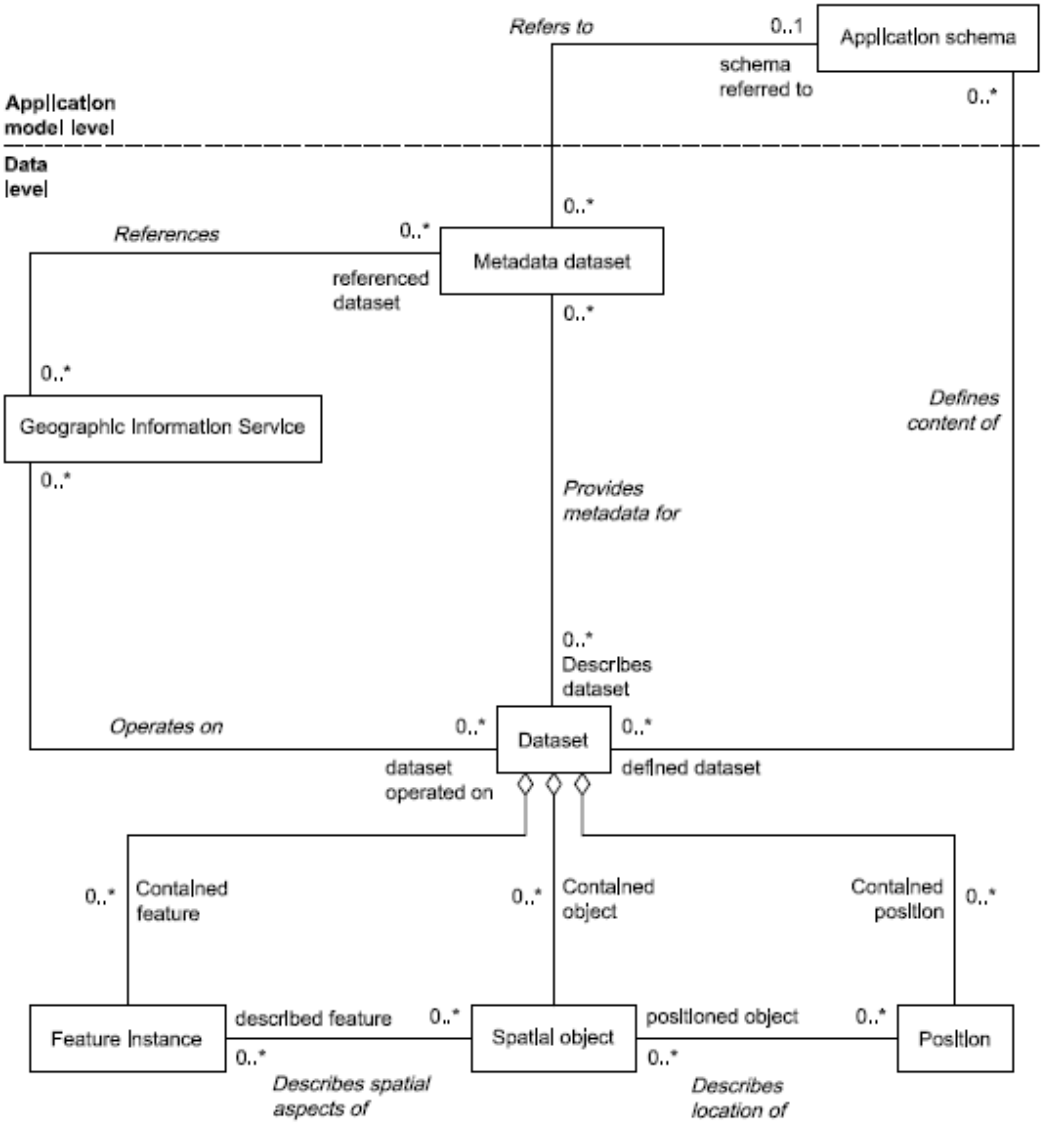

Figure A-2: High-level view of the Domain reference model (ISO 19101 Figure 5)

Figure Figure provides a high-level summary of standardization requirements.

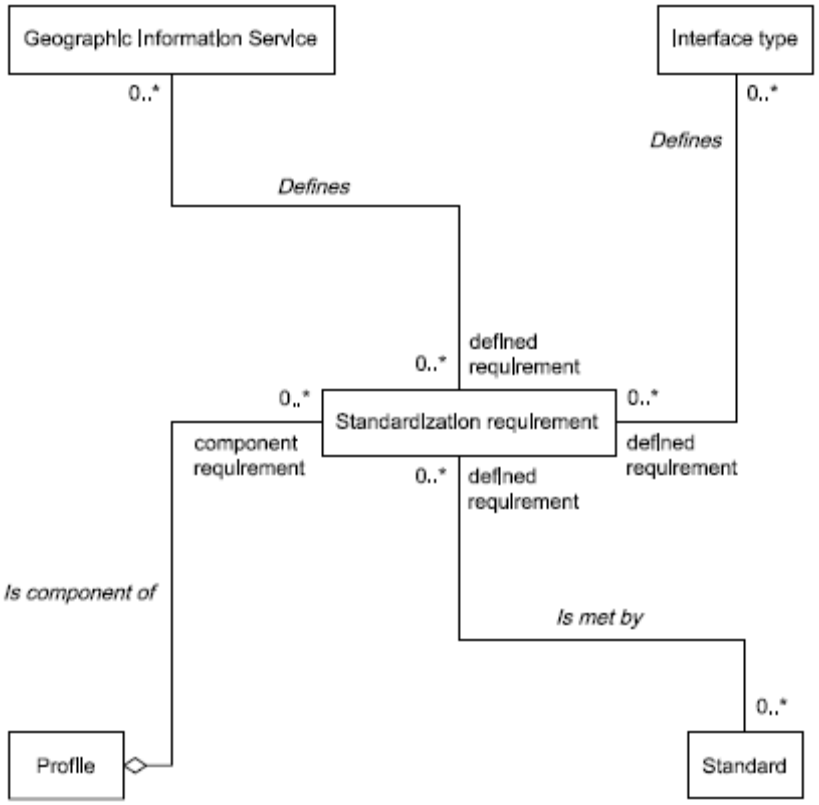

Figure A-1: Identification of standardization requirements (ISO 19101 Figure 14)

Two other ISO documents establish the modelling framework and notation, including a high-level meta-model. These are:

1. ISO TS 19103:2005 Geographic information—Conceptual schema language – which provides:
    a. a profile of UML to be used in domain modelling, focussing particularly on class diagrams; and
    b. a set of primitive data-types, covering numbers, text, measures, collections, and records.
2. ISO 19109:2005 Geographic information—Rules for application schema – which defines a hierarchy of meta-levels for modelling application domains (such as Water Resources), based on the General Feature Model (GFM).

In the remainder of this Appendix, the significant components provided by each of the ISO standards are summarised. Note that the intention of these standards is to provide a framework for information modelling and service architectures for GI. They are mostly expressed as UML models, and thus software implementation requires the application of platform specific rules on top of the models. In particular any domain model is constructed by combination and sometimes specialization of the elements provided.

A.2. Notation - ISO 19103

A profile of UML, including recommendations concerning the avoidance of multiple inheritance, labelling of associations with role-names, and some standard class stereotypes.

A.3. Feature Model, Feature Type Catalogue-GFM ISO 19109, ISO 19110

ISO 19109 introduces a key principle of the ISO framework, that an information model and data transfer is tied to a particular domain-of-discourse or community. The information model embodies the concepts, including data-types, which characterise discourse within a particular domain. Membership of the domain is indicated by use of this application schema or language.

The General Feature Model (GFM) is a meta-model for the “fundamental unit of geographic information” known as a feature. Features may correspond to any item of interest, not restricted to tangible real-world objects, although it is common for feature instances to correspond with conventional real-world objects (e.g. dam-A, stream-B, catchment-C). The feature model is thus essentially an object model, with the key requirement that each feature instance is typed and identifiable. Feature type is defined by a characteristic set of properties, which may include attributes, associations, and operations.

A feature-type-catalogue (FTC) is the set of feature-types, which effectively comprise the nouns of a domain of discourse or application schema.

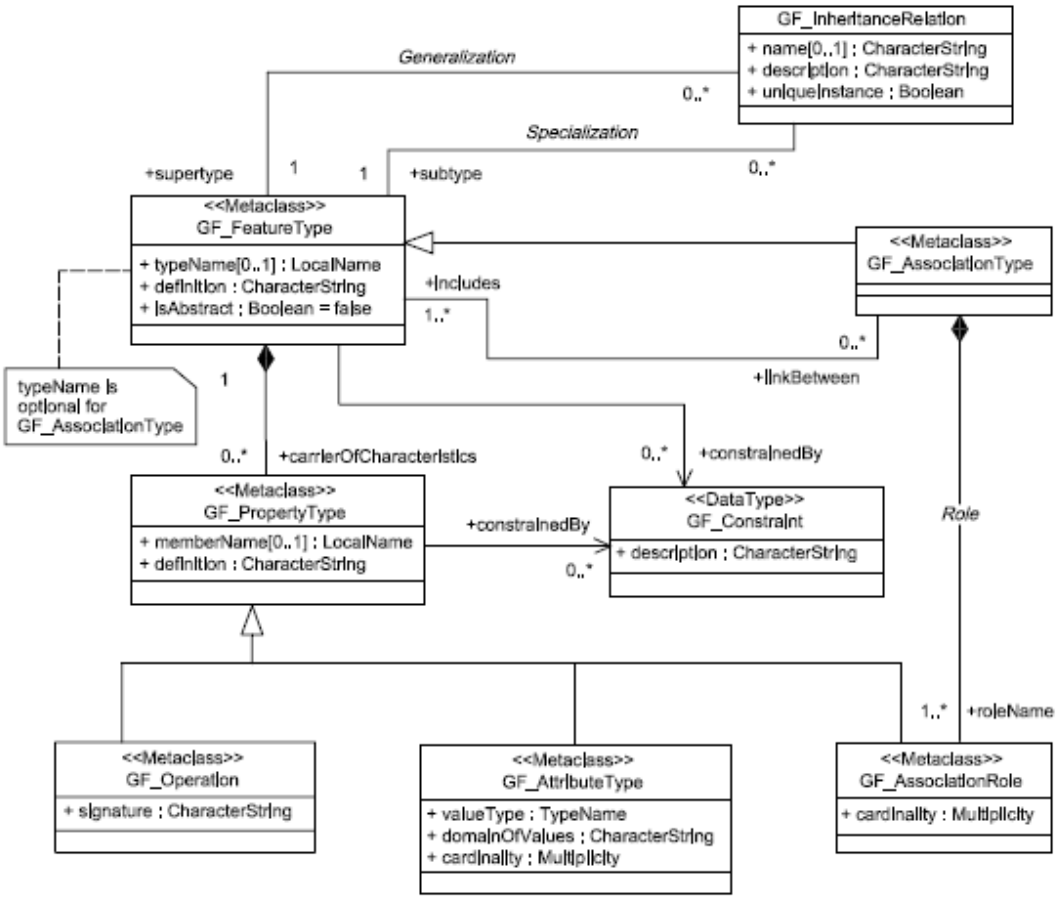

Figure A-2: Extract from the General Feature Model [ISO 19109 Figure 5]

A.4. Objects and Coverages- ISO 19109, 19123

The GFM defines the information units that are of interest for discourse and thus transfer. Many properties of a feature instance are constant on the feature (e.g. name, owner, shape) so characterisation of the world may be achieved by considering feature instances and their static

properties. However, for some applications the variation of a property in a spatio-temporal domain is the primary focus. This viewpoint is termed a coverage, which maps the values in a property range onto a domain (the 'independent variable') composed of spatio-temporal elements. The coverage domain will correspond with the extent of some feature instance, of a type that may or may not be of direct interest, such as a *scene*. While a property may vary continuously within the coverage domain, a coverage is usually sampled at discrete locations within the domain, such as grid points. A function may be available to define values in the range for locations that are not sampled directly.

For the WRON, all monitoring applications can be represented as a coverage, to express the variation of the monitored property with time.

A.5. Utility Types- ISO 19103, 19107, 19108, 19111

Definitions of sets of important generic data types for GI are the subject of several of the ISO 19100 standards. These are all expressed as UML class models.

19103 provides basic data-types, for numbers, text, truth, measures, aggregates, and records.

19107 provides a comprehensive spatial schema, including definitions for the key geometry types: points, various kinds of curves, surfaces and solids. A topology model is also defined.

19108 provides a schema for temporal objects (including temporal topology) and reference systems (including ordinal reference systems).

19111 provides a schema for Coordinate Reference Systems, and its building blocks including reference systems, coordinate systems, datums and coordinate transformations.

A.6. Metadata

Version 0.1: To be completed in version 1.0 (Summarize scope of ISO 19115)

A.7. Registry Model

Version 0.1: To be completed in version 1.0 (Summarize ISO 19135)

A.8. Profiles

Version 0.1: To be completed in version 1.0 (Summarize ISO 19106)

A.9. XML Encoding

Version 0.1: To be completed in version

B. ISO 19100 Framework

Observations and Measurements (O&M) defines a generic model for an observation event, which estimates the value of a feature-property (**Error! Reference source not found.**). The Observation type provides access to metadata about the observation event, which is of interest when evaluation of data quality is of interest. For a generic observation, the feature-of-interest is 'any feature', and the other properties of the Observation type provide the parameters for a request for observations – e.g. "give me an observation of the temperature of John Doe, sampled yesterday using a mercury-in-glass-thermometer". Observation results may be scalar or complex data structures, including coverages such as time-series.

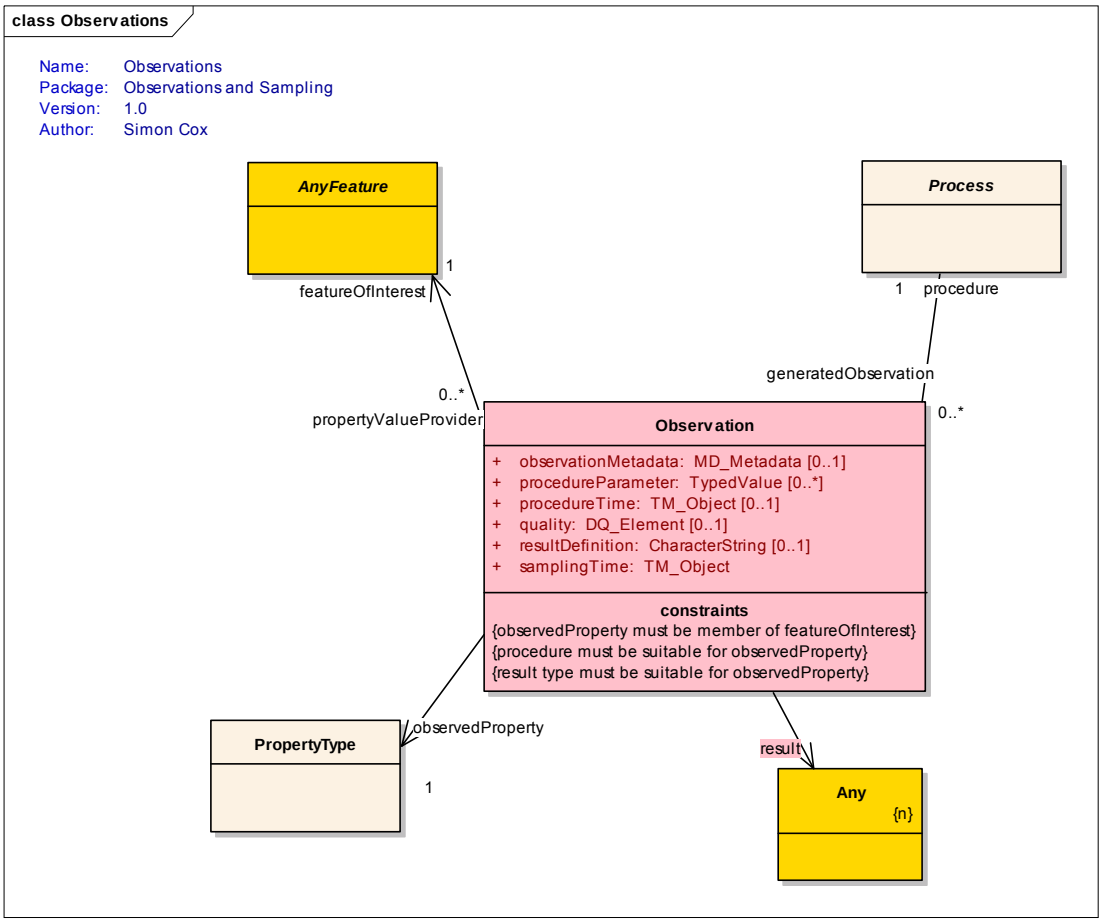

Figure B-1: Observations and Measurements Pattern.

Sampling features embody strategies commonly used to make observations in the natural environment, when exhaustive observation of a feature instance is not possible. For example, an aquifer is commonly sampled in a well; an atmosphere or ocean column on a sounding; the weather is sampled at stations; rock-units are observed in cross-sections and outcrops. Sampling Features provides a general model for this and the pattern that unifies this with the Observations and Measurement pattern can be seen below in Figure B-2: Sampling Feature related to Observations Pattern..

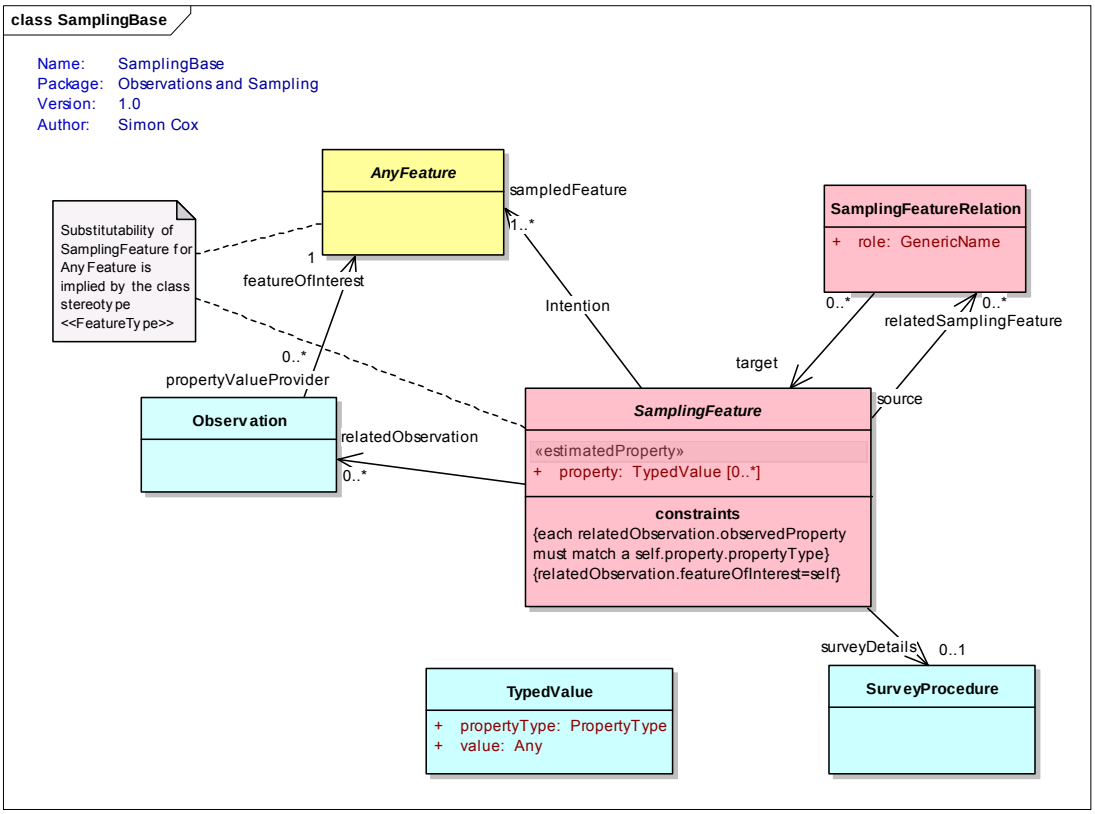

Figure B-2: Sampling Feature related to Observations Pattern.

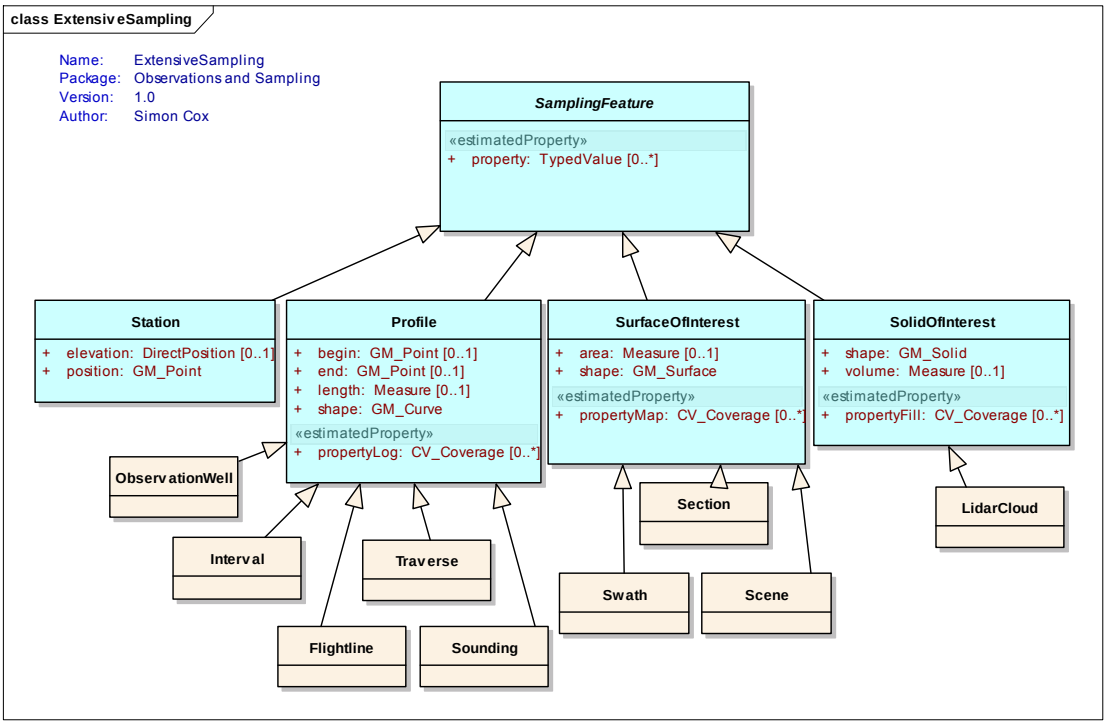

Figure B-3: Extensive Sampling Pattern

10. Bibliography

Budd, K., Hostetler, S., Gilmour, J., Chiles, G., Bleys, E., and Brodie, R., 2004, 'Review of the Australian Water Resources Assessment 2000', AWDIP Report 2.

CSIRO, 2006, 'Water for a Healthy Country Water Resources Observation Network (WRON) Theme 1010 Science Investment Process Plan', Internal Report (SIP2).

Dijkstra, E. W., 1982 'Selected Writings on Computing: A Personal Perspective', Springer-Verlag, pp 60-66, 1982.

Elkin, P., Bauer, B., Wahner-Roedler, D., Rosenbloom, S., Bailey, K., Froehling, D., Brown, S., 'Aequus Communis Sententia: Defining Levels of Interoperability', MEDinfo 2007 accepted paper, IOS Press, 2007.

Fitzke, J., Atkinson, A., 2006, 'Gazetteer Service – Application Profile of the Web Feature Service Implementation', OGC Best Practices Document OGC 05-035r2.

Foster, I., 2006, 'Globus Toolkit Version 4: Software for Service-Oriented Systems', IFIP International Conference on Network and Parallel Computing, Springer-Verlag LNCS 3779, pp 2-13, 2006.

Gruber, T., 1993, 'Towards Principles for the Design of Ontologies Used for Knowledge Sharing', International Journal Human-Computer Studies No. 43, pp. 907-928.

Hasselmeyer, P., 2005, 'On Service Discovery Process Types', 3rd International Conference on Service Oriented Computing, Amsterdam, Springer-Verlag LNCS 3826, pp 144-157, 2005.

National Water Commission, 2005, Objectives of AWR 005, Australian Water Resources 2005 website (http://www.water.gov.au/IntroductiontoAWR2005/ObjectivesofAWR2005/index.aspx?Menu=Level1_2_2).

National Water Commission, 2006, 'Australian Water Resources 2005 – A baseline assessment of water resources for the National Water Initiative, Level 1 Assessment', National Water Commission, Canberra.

Renaud, G., 2004, 'Metadata and Controlled Vocabularies in the government of Canada: A Situational Analysis', International Conference on Dublin Core and Metadata Applications, Shanghai, 2004.

van Dijk, A. et.al, 2006, 'Risks to the Shared Water Resources of the Murray-Darling Basin', Murray-Darling Basin Commission, Canberra.

Water Services Association of Australia, 2005, 'Testing the Water – Urban water in our growing cities: the risks, challenges, innovation and planning', WSAA Position Paper No. 01.