

Repository Federation Service Usage Model: provision and manage repository federation – federated metadata registry

Repository Federation Service Usage Model: provision and manage repository federation – federated metadata registry	1
1 Introduction.....	2
2 Service Usage Model Definition	3
2.1 Rationale	3
2.2 Service Usage Model Metadata.....	3
2.3 Name	3
2.4 Classification	3
2.5 Version.....	4
2.6 Notation	5
2.7 Description	6
2.8 Business Process Modelling	6
2.9 Usage Scenarios	9
2.10 Applicability.....	9
2.11 Functionality.....	10
2.12 Structure & Organization.....	15
2.13 Design Decisions & Tradeoffs.....	22
2.14 Implementation Guide & Dependencies.....	23
2.15 Applicable Standards	24
2.16 Known Uses	24
2.17 Resources Used	25
2.18 Services (Service Genres) Used	25
2.19 Related Service Usage Models.....	29
2.20 Service Patterns/CORE Service Usage Models Used	30
2.21 References	31

Copyright © University of Southern Queensland and University of Memphis.



This work is licensed under the Creative Commons Attribution-Share Alike 2.5 Australia License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-sa/2.5/au/>

This work is created as part of the Federated Repositories for Education (FRED) Project within the Australian ADL Partnership Laboratory. The FRED project is sponsored by the Australian Commonwealth Department of Education, Science and Training under the Framework for Open Learning Programme.

2.22 Glossary	31
---------------------	----

1 Introduction

The Introduction provides a brief, standalone overview of the Service Usage Model. It is for a non technical reader. It is not for inclusion in a submission to the e-Framework, and thus may duplicate other material in the Service Usage Model Definition.

Repository federation is the process by which a set of repositories or content collections are combined into a single (logical) entity to support the discovery of the content contained within the individual repositories that form the federation. For the end user, the federation (as a software system) provides the single access point to the content in the federation, eliminating the need for a user to identify and individually access (search) each of the repositories to find content. The end user goes to the federation service end point to discover and identify content and its location (in the constituent repositories in the repository federation). The user may then go to an identified source repository to obtain the identified content.

The specific assumptions and features of such a repository federation, and its functionality exposed to end users, are described in the repository federation (content) service usage model. That service usage model was written in the FRED project, to model such repository federations. This service usage model provides provisioning and management functionality for the type of repository federation modelled in the repository federation (content) service usage model. In general, the provisioning and management functionality is independent of the details of the repository federation. Specific core assumptions for the federated metadata registry provisioning and management service usage model description are:

- The repository federation includes a central registry (a resource) that provides the service end points for all services exposed by the service usage model.
- Access to the services of the repository federation may require authorization, and authentication processes may be applied to determine access conditions and to limit results.
- Access to the services of the participating repositories may require authorization, and authentication processes may be applied to determine access conditions.

The overall model of the repository federation is illustrated in Figure 1.

The service usage model covers the provisioning of a repository federation (e.g., provisioning policies, business rules, users), managing discovery of content objects (e.g., computing ranking, ratings, recommendations), or central registry operations (e.g., backup).

The services exposed by the federated metadata registry provisioning and management service usage model may be used to build end-user applications, e.g., a content management portal; or the services may be integrated into other applications, e.g., integrated with identifier and identity management.

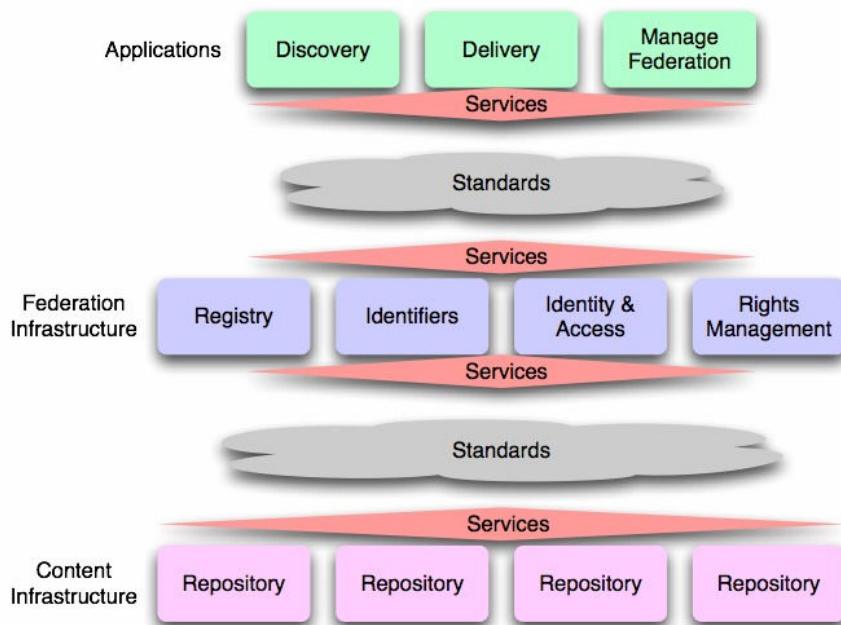


Figure 1: Repository Federation Model

2 Service Usage Model Definition

The Service Usage Model Definition is the complete formal documentation of the Service Usage Model for submission to the e-Framework. The elements included extend those from the current e-Framework template.

2.1 Rationale

The Rationale element presents the justification for the Service Usage Model, and explains its use intended by its authors.

The FRED project has specified, in the repository federation (content) Service Usage Model, the range of services expected to be used in repository federations in education in Australia for interacting with content objects (namely the functionality relating to objects, as distinct from the overall repository federation systems). This Service Usage Model outlines the range of registry provision and management services expected to be used in supporting such a repository federation.

2.2 Service Usage Model Metadata

The Service Usage Model Metadata contains basic labelling, classification and a version history for the Service Usage Model. It is expected that the e-Framework will modify these elements after submission.

2.3 Name

- FRED Service Usage Model Name: Federated Metadata Registry Provisioning and Management
- e-Framework Service Usage Model Registry Name: TBD

2.4 Classification

- Type: Domain

- Status: Unapproved
- Domains: Library (some of the system management functions might be appropriate for the IT domain)
- Domain Coverage: Single
- Deployment: Isolated
- Development Status: Proposed
- XOR: Genre
- Maturity: Immature
- Composition: Composite
- Purpose: Exemplar
- End Point: Transcoder
- Auth'd: Auth'd

2.5 Version

- FRED Version: 1.0.2
- e-Framework Service Usage Model Version: TBD

Version History			
Version	Date	Author	Description
0.1	2007-02-14	DR	Initial partial version. hdl:FREDNA/4C6E74F3FEDF4784B6CED2F08D43F95D
0.2	2007-02-21	DR	Partial version, most functionality and services.
0.3	2007-02-25	DR	Complete draft. See "to do" list at the end of the document. Unedited
0.31	2007-02-26	DR	Complete draft. See "to do" list at the end of the document. Edited
0.32	2007-03-14	NN	Queries, minor editorial
0.4	2007-03-26	DR	Edits, new diagrams (4, 5), modified some service descriptions, query responses, embedded Visio diagrams (1,2,4,5,6, not 3)
0.45	2007-04-10	DR	Added element descriptions, shuffled sections, editorial, accepted some changes from V0.4
0.46	2007-04-13	NN	Added a rationale and business analysis component
0.47	2007-04-25	NN	Editorial
0.48	2007-04-	NN	Editorial: global change of service + resource nomenclature to service {resource}; conflate query

	27		functionalities; align diagram of business processes to rest of business analysis
0.49	2007-05-07	DR	Editorial, Accepted lots of changes. Deleted notes that have been resolved.
0.5	2007-05-09	NN	Editorial, tidying up outstanding comments, proceeded with conflation of similar functionalities.
0.51	2007-05-14	DR	More editorial and cleanup
0.52	2007-05-24	NN	Cleanup: Service Genre listings has genres not genre instances; notation names data source classes not instances; functionalities unconfliated; replaced core/value-added with base/optional; moved consumer SUMs to Known Uses and out of Related SUMs.
0.53	2007-05-25	DR	Editorial. Revise Genre listing. Changes to get ready to submit to the e-Framework.
1.0.1	2008-06-18	DR	Editorial to align with LSAL ADL work.
1.0.2	2008-07-17	NN	Editorial

Notation

The Notation element includes conventions used to describe the Service Usage Model.

The words MUST, MUST NOT, REQUIRED, SHALL, SHALL NOT, SHOULD, SHOULD NOT, RECOMMENDED, MAY, and OPTIONAL in this document are to be interpreted as described in [RFC 2119].

Definitions of service usage model-specific terms used throughout (e.g., content object, collection, collection object, metadata object, registry, registry object, repository, repository object) are given in the service usage model glossary. Other terms, e.g., client, provider, resource, are used as defined in the e-Framework glossary.

The service usage model description that follows uses e-Framework service genre description elements as of 2006-12-09, updated to include draft e-Framework service classifications (V0.5) and revised to incorporate other elements based on FRED project experience needed to document the service usage model.

Service genres are named using a proposed e-Framework naming convention, consisting of the behaviour or function provided by the service genre (expressed as a verb) augmented with the name of the target resource class (or object class) in curly brackets; e.g., authenticate {agent} (applies to any agent), obtain {content repository} (applies to any content repository in the federation), validate {metadata object} (applies to any metadata object) . Service genres and resources are also classified by type, e.g., BASE (essential), VALUE ADDED (not essential), EXTERNAL (exposed to users, default if not noted), or INTERNAL (not exposed to users).

Service usage models are classified by type, e.g., BASE (essential), VALUE ADDED (not essential). Related service usage model that describes an Application which relies on the service usage model are denotes APPLICATION. If not classified, the related service usage model may be use by, or may use this service usage model.

Items are tagged and identified in the metadata element using names assigned by the FRED project. Formal e-Framework names will be assigned by the e-Framework.

Prior to e-Framework publication, the version numbers shall be less than 1.0 and a new identifier is not assigned to each version. Revisions subsequent to publication by the e-Framework require a new identifier be assigned to each version.

2.7 Description

The Description element is an informal, standalone, non technical narrative description of the Service Usage Model (problem, process, business level capabilities and workflow).

The federated metadata registry provisioning and management service usage model describes the structure of a collection of collaborating service genres and service usage models (along with their associated managed resources) used to provide all the functionality needed to provision and manage a repository federation and metadata registry.

The repository federation is an organized collection of constituent repositories and a central registry used to discover content from the repositories that participate in the repository federation. The service genres within the service usage model provide the functionality to provision the resources that support the capabilities of the federation and to manage the central registry.

The federated metadata registry provisioning and management service usage model provides capabilities to:

- Provision the key resources used in the operating the repository federation.
- Manage the central registry and its constituent parts.

The basic structure of the federated metadata registry provisioning and management service usage model is a collection of collaborating service genres operating on the central registry resources. The service usage model defines a set of business functions that can be exposed to users; each function maps to a service end point for the repository federation. Each service end point responds to a user request and either adds information to the central registry or manipulates information that is part of the central registry. Individual and combined service genres, with a prescribed workflow, provide the described functionality.

The overall business functionality of the federated metadata registry provisioning and management service usage model is illustrated in Figure 2.

The federated metadata registry provisioning and management service usage model is used with other service usage models, including identity management service usage models and identifier service usage models. The federated metadata registry provisioning and management service usage model is used in conjunction with the repository federation (content) Service Usage Model to provide complete functionality for repository federation.

The defined federated metadata registry provisioning and management service usage model must be further refined for a community of practice. This specialization will add policy, defined data models, standards used and behaviour constraints. These will be used to refine the service genres into service expressions and aid in developing the design for the complete application suite needed to provision and manage a repository federation and metadata registry.

2.8 Business Process Modelling

The Business Process Modelling element lists the business functions to be supported by the Service Usage Model. Business process modelling provides necessary context to motivate and understand the

description in the remainder of the Service Usage Model. The business analysis motivating the Service Usage Model should be summarised within the Service Usage Model (to make it readable as a self-contained object), with fuller discussion possibly outside the Service Usage Model. Functions should be listed as groups, in business terms. The description should not be encumbered with either technical services or detailed workflows.

Under *Functionality*, services are split into Registry Domain functions, performed by registry users, and Registry System functions, used to manage the repository federation. Registry Domain functions are not discussed in this service usage model, and are covered in the repository federation (content) Service Usage Model.

The following Registry System functions have been identified and are supported by services in the Service Usage Model. Note that detailed business process modelling has not been completed for this Service Usage Model, and there are no high-level business functions defined to correspond to these system functions. The functionality of this Service Usage Model is driven by the operation requirements of the repository federation (content) Service Usage Model, and do not have an independent business motivation. For example, *Compute Discovery Ranking* provides operational support for the content SUM function *Query Content Object Registry* in the Repository Federation (content) Service Usage Model, which in turn is motivated by the business function *Discover Item through Search on Registry*.

Registry Provisioning:

- *Provision Authorization Policies*
- *Provision Authentication*
- *Provision Users*
- *Provision Business Rules*. These control registry operations, and are used in deciding whether various requests made of the repository federation should be fulfilled, and how; e.g. authorizing requests to update metadata objects, or to query content objects.
- *Provision System Data*. These are system-level objects used to describe the registry and its content, such as vocabularies and schemata.

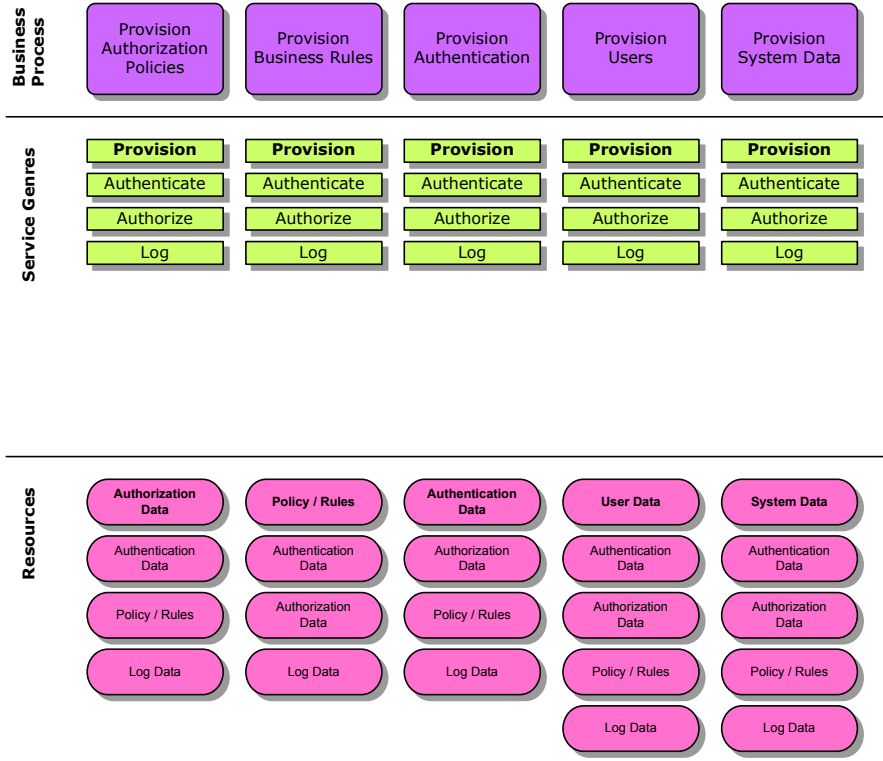
Registry Management:

- *Adjust Registry Search Index*
- *Compute Discovery Ranking*
- *Compute Discovery Classification*
- *Determine Discovery Recommendations*
- *Compute Relationships*
- *Monitor Registry Behaviour*
- *Inspect Registry*
- *Analyze Registry*

NB: The Service Usage Model may also be used to support additional high-level business functions that are outside of the scope of this analysis and FRED.

Service Usage Model Diagram

Service Usage Model diagrams showing the business functions, service genres and resources are shown in Figures 2 and 3.

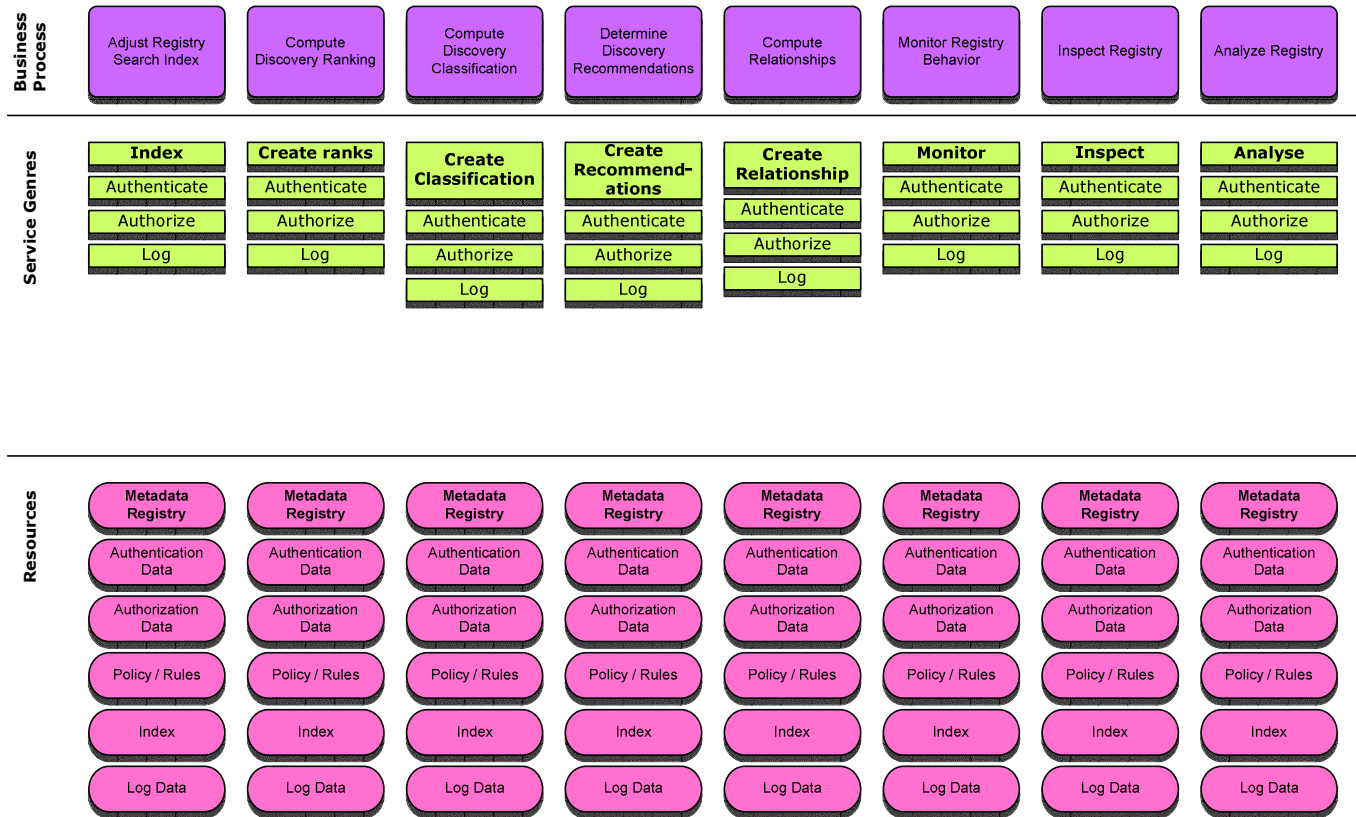


© Copyright 2007, University of Southern Queensland and University of Memphis

This work is licensed under the Creative Commons Attribution-ShareAlike 2.5 Australia License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-sa/2.5/au/>



Figure 2: Service Usage Model Diagram: Registry Provisioning



© Copyright 2007, University of Southern Queensland and University of Memphis

This work is licensed under the Creative Commons Attribution-ShareAlike 2.5 Australia License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-sa/2.5/au/>



Figure 3: Service Usage Model Diagram: Registry Management

2.9 Usage Scenarios

The Usage Scenarios element is an informal, non technical description of how the Service Usage Model is used in support of business processes. An illustration of business workflows, expressed using services, is included. An illustration of an application using the components of the Service Usage Model may be included (but not a description of the design of the application). Full business process modelling and detailed diagrams are excluded. No critical or essential information required to understand the Service Usage Model should be included.

Provision and management usage scenarios for the (FRED) repository federation are documented separately in FRED documentation [[reference to FRED use cases](#)]

2.10 Applicability

The Applicability element details how and when the Service Usage Model is used or not used. It represents specific constraints and assumptions on the use of the Service Usage Model. It is more specific and normative than the informal Usage Scenarios. No critical or essential information required to understand the Service Usage Model should be included.

The federated metadata registry provisioning and management service usage model focuses on provision and management of the metadata registry, independent of how the metadata registry is used to create a repository federation. It presupposes the metadata registry operates as specified in the repository federation (content) Service Usage Model.

The federated metadata registry provisioning and management service usage model defines a trust environment. External access to the services and resources that are within the trust environment requires authorization or access controls. Within the trust environment, access to individual services and resources from other services within the trust environment SHOULD NOT require authorization or access controls.

The federated metadata registry provisioning and management service usage model is based on a common identifier model and system used throughout to identify all objects and provide resolution services.

The federated metadata registry provisioning and management service usage model is applicable both for creating end-user provisioning and management applications and imbedding provisioning and management functionality in other applications and tools.

The federated metadata registry provisioning and management service usage model does not specify functional requirements for portals, human user interfaces or capabilities (customization, personalization, help, etc.) used to develop end-user applications.

The federated metadata registry provisioning and management service usage model does not specify security requirements..

2.11 Functionality

The Functionality element details and illustrates the behaviours provided by the Service Usage Model, in terms of services, workflows, messages, resources, data objects. It is not a technical description of the Service Usage Model, but it must provide sufficient information to develop the Structure & Organization of the Service Usage Model and to evaluate conformance of the Service Usage Model to the stated behaviours. It should not include implementation-specific information.

The functionality required in the deployment of a repository federation can be split into two major groups:

- *Registry Domain Functions*: functions performed by users (end users, services) dealing with content objects in repositories and registries, including building federations and the discovery and retrieval of content objects from federations and their constituent repositories.
- *Registry System Functions*: functions applied to the registry infrastructure management and operations. These are further subdivided as:
 - *Registry Provisioning Functions*: Creating a registry (and its constituent parts) to enable building a repository federation.
 - *Registry Management Functions*: Managing the operations of a registry.

This service usage model only covers *Registry System Functions*. The *Registry Domain Functions* are described in the repository federation (content) Service Usage Model.

These functions are further classified as:

- BASE (essential capabilities) versus VALUE ADDED (desirable but not essential repository federation services). The distinction is based both on technical requirements and community requirements.

NB: The functionality described does not imply a particular implementation strategy.

2.11.1 Registry Provisioning Functions

Provision Authorization Policies: [VALUE ADDED]

Establish or modify any of the authorization policies used to control access to the components and resources of the repository federation.

Basic workflow steps (without error handling) are:

- Provide the authorization policies to be provisioned (the request).
- Authenticate and authorise (via access controls and business rules) the request.
- Validate the request and policies.
- Update the access control and authorization policies.
- Log the activity.
- Return status results.

Business Process: —

Provision Business Rules: [VALUE ADDED]

Establish or modify any of the business policies and business rules used to control the behaviour of the components of the repository federation.

Basic workflow steps (without error handling) are:

- Provide the business rules and policies to be provisioned (the request).
- Authenticate and authorise (via access controls and business rules) the request.
- Validate the request and business rules.
- Update the business rules and policies.
- Log the activity.
- Return status results.

Business Process: —

Provision Authentication: [VALUE ADDED]

Establish or modify any of the authentication data used to authenticate access to the components of the repository federation.

Basic workflow steps (without error handling) are:

- Provide the authentication policies to be provisioned (the request).
- Authenticate and authorise (via access controls and business rules) the request.
- Validate the request and policies.
- Update the authentication policies.
- Log the activity.
- Return status results.

Business Process: —

Provision Users: [VALUE ADDED]

Establish or modify the collection of users, their roles and trust boundaries for the components of the repository federation.

Basic workflow steps (without error handling) are:

- Provide the user identity information to be provisioned (the request).
- Authenticate and authorise (via access controls and business rules) the request.
- Validate the request and business rules.
- Update user identity information (e.g., user identity data, user authentication data).
- Log the activity.
- Return status results.

Business Process: —

Provision System Data: [VALUE ADDED]

Establish or modify any of the system data structures (schemata, vocabularies) for the repository federation.

Basic workflow steps (without error handling) are:

- Provide the system data to be provisioned (the request).
- Authenticate and authorise (via access controls and business rules) the request.
- Validate the request and business rules.
- Update the system data.
- Log the activity.
- Return status results.

Business Process: —

2.11.2 Registry Management Functions

Adjust Registry Search Index: [BASE]

Modify indexing parameters and reindex the registry.

Basic workflow steps (without error handling) are:

- Provide a new set of indexing parameters (the request).
- Authenticate and authorise (via access controls and business rules) the request.
- Validate the request.
- Update the indexing parameters in the system data.
- Reindex the registry.
- Log the activity.
- Return status results.

Business Process: —

Compute Content Object Discovery Ranking: [VALUE ADDED]

Compute ranking metadata for all discoverable content in the repository federation through an analysis of data available in the registry, i.e., compute and add ranking metadata and descriptive information to a metadata object in the metadata registry.

Basic workflow steps (without error handling) are:

- Provide a 'compute ranking' request (the request).
- Authenticate and authorise (via access controls and business rules) the request.
- Validate the request.
- Compute the rankings. Update the catalogue of metadata objects in the repository federation.
- Log the activity.
- Return status results.

Business Process: —

Compute Content Object Discovery Rating: [VALUE ADDED]

Compute rating metadata for all discoverable content in the repository federation through an analysis of data available in the registry, i.e., compute and add rating metadata and descriptive information to a metadata object in the metadata registry.

Basic workflow steps (without error handling) are:

- Provide a 'compute rating' request (the request).
- Authenticate and authorise (via access controls and business rules) the request.
- Validate the request.
- Compute the ratings. Update the catalogue of metadata objects in the repository federation.
- Log the activity.
- Return status results.

Business Process: —

Compute Content Object Discovery Classifications: [VALUE ADDED]

Compute classification metadata for all discoverable content in the repository federation through an analysis of data available in the registry, i.e., compute and add classification metadata and descriptive information to a metadata object in the metadata registry.

Basic workflow steps (without error handling) are:

- Provide a 'compute classification' request (the request).
- Authenticate and authorise (via access controls and business rules) the request.
- Validate the request.
- Compute the classifications. Update the catalogue of metadata objects in the repository federation.
- Log the activity.
- Return status results.

Business Process: —

Determine Metadata Object Discovery Recommendations: [VALUE ADDED]

Determine recommendations for alternative content objects for all discoverable content in the repository federation through an analysis of data available in the registry, i.e., compute

and add recommendation metadata and descriptive information to a metadata object in the metadata registry.

Basic workflow steps (without error handling) are:

- Provide a 'compute recommendation' request (the request).
- Authenticate and authorise (via access controls and business rules) the request.
- Validate the request.
- Compute the recommendations. Update the catalogue of metadata objects in the repository federation.
- Log the activity.
- Return status results.

Business Process: —

Compute Relationships: [VALUE ADDED]

Determine FRBR relationships (expressions, manifestations, copies) between content objects for all discoverable content in the repository federation through an analysis of data available in the registry, i.e., compute and add relationship links, relationship metadata and descriptive information to a metadata object in the metadata registry.

Basic workflow steps (without error handling) are:

- Provide a 'compute relationship' request (the request).
- Authenticate and authorise (via access controls and business rules) the request.
- Validate the request.
- Compute the relationships. Update the catalogue of metadata objects in the repository federation. Update the relationship data.
- Log the activity.
- Return status results.

Business Process: —

Monitor Registry Behaviour: [VALUE ADDED]

Monitor the processes and overall behaviour of the registry and its component systems.

Basic workflow steps (without error handling) are:

- Provide a request for component (hardware, software system) status (the request).
- Authenticate and authorise (via access controls and business rules) the request.
- Validate the request.
- Probe the component to determine its status.
- Return status results.

Business Process: —

Inspect Registry: [VALUE ADDED]

Inspect the contents of the registry and its component systems.

Basic workflow steps (without error handling) are:

- Provide a request for component (hardware, system) status (the request).

- Authenticate and authorise (via access controls and business rules) the request.
- Validate the request.
- Inspect the registry contents.
- Return status results.

Business Process: —

Analyse Registry: [VALUE ADDED]

Provide analytical data on registry performance and usage.

Basic workflow steps (without error handling) are:

- Provide a request for registry analytics (the request). The request will specify the type of analytics data and dissemination format.
- Authenticate and authorise (via access controls and business rules) the request.
- Validate the request.
- Retrieve analytics from the registry logs.
- Filter the analytics results (via content access controls).
- Transform the analytics results into the requested format.
- Log the activity.
- Return the analytics results.
- Return status results.

Business Process: —

Backup Registry: [BASE]

Provide operational registry storage component backup.

Basic workflow steps (without error handling) are:

- Provide a backup request (the request). The request will specify which registry storage component to backup.
- Authenticate and authorise (via access controls and business rules) the request.
- Validate the request.
- Perform the backup.
- Log the activity.
- Return status results.

Business Process: —

Other Functions

No other functionality is defined. The functionality that is defined MAY be extended.

2.12 Structure & Organization

The Structure & Organization element is a normative, technical description of the Service Usage Model that documents and illustrates the service-oriented operational view of the Service Usage Model, in terms of services, resources and their messaging interactions. The Structure & Organization element is required to implement applications that use the Service Usage Model, but the Service Usage Model must be understandable without reference to the contents of the element.

NB: The Structure & Organization section does not describe how to implement applications that use the Service Usage Model, but describes how to implement elements of the Service Usage Model itself.

The structure of the repository federation is illustrated in Figure 1. The core of the repository federation is the central registry, which includes management resources such as system data, policy and rule data, and logs. The central registry also uses resources that are managed by other service usage models, including access control, user and identity data that are key resources of an identity service usage model and identifier resolution data that is part of an identifier service usage model.

A revised view of the federated metadata registry provisioning and management service usage model, combined with Domain System services (repository federation (content) Service Usage Model) to provide complete repository federation functionality, is illustrated in Figure 4. Blue dots on services indicate that the service is applied to the central registry resources (e.g., metadata registry) (also marked with blue dots). The diagram includes *System* services in the federated metadata registry provisioning and management service usage model and *Domain* services in the repository federation (content) Service Usage Model. It does not illustrate the identity provisioning and management or identifier provisioning and management service usage models.

The federated metadata registry provisioning and management service usage model exposes a collection of service end points, modelled as service genres. The core functionality of the service usage model is mapped to these service genres. Typically one function is mapped to one service end point, supported by a single service genre. The service genres function in the same fashion, accepting a request from the user and then performing the requested operation. When valid, the operation either modifies the resources of the central registry or updates information for the objects in the central registry. The workflows for all functions follow this similar pattern.

Other than working on a shared resource set, the service genres are not interdependent. All service end points do rely on a common set of service genres, including those for identity management and access control.

Most of the service genres function within a trust environment. External requests are validated, but inside the trust environment, service requests are accepted as is from trusted services.

Figure 5 illustrates an Enterprise Service Bus-oriented (ESB) view of a typical workflow in the federated metadata registry provisioning and management service usage model. The shared resources are shown at the bottom on the diagram. The identifier and identity resources are assumed to be part of an external environment.

There is no defined overall application. The federated metadata registry provisioning and management service usage model is assumed to be used within a set of applications that are used to provide provisioning and management of the repository federation, provide management and provide operational support.

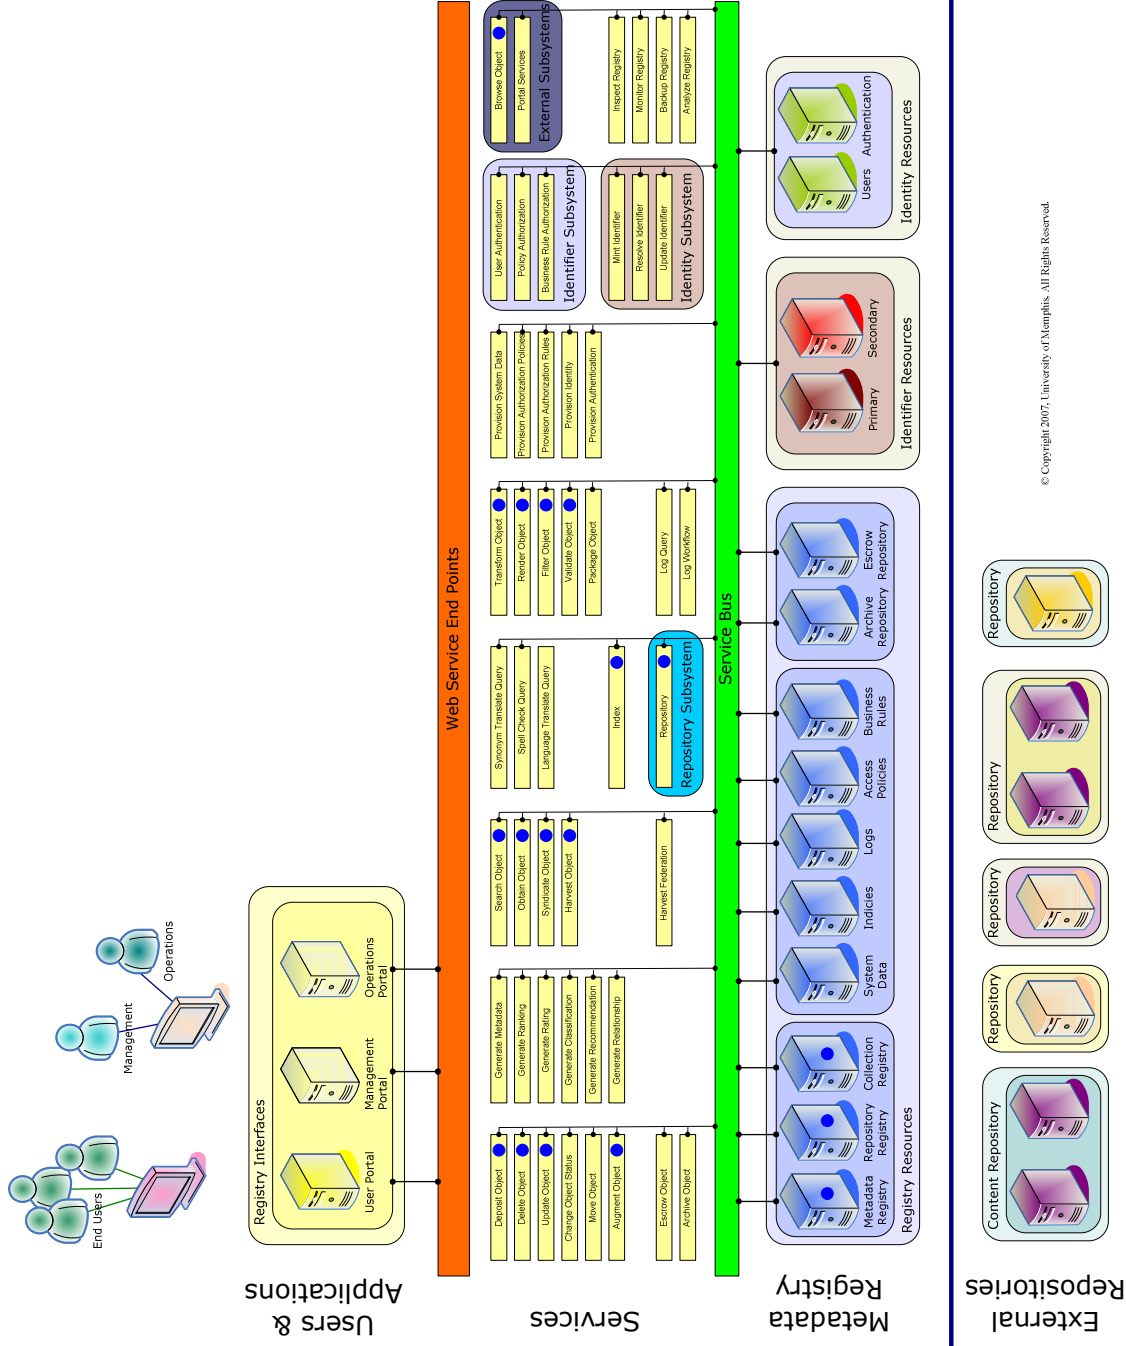


Figure 5: Repository Federation Service Usage Model

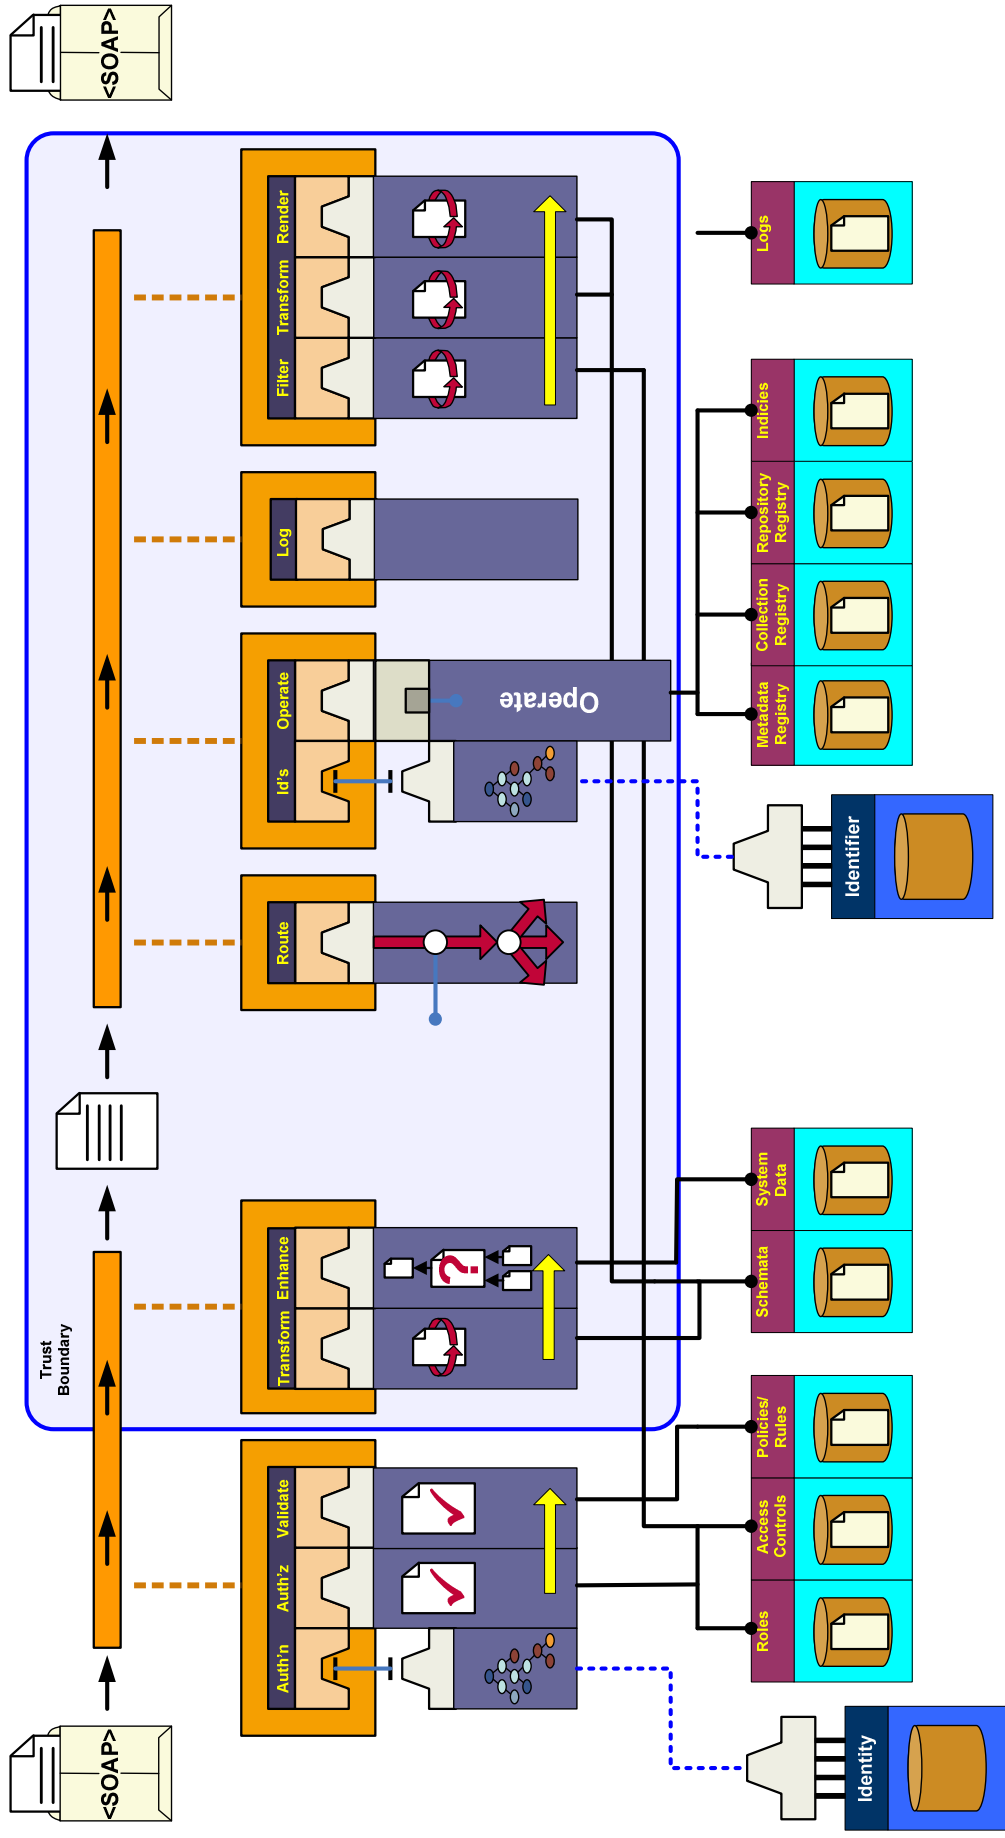


Figure 6: Repository Federation Service Usage Model (ESB Notation)

The functionality outlined above is realised through services as follows:

2.12.1 Registry Provisioning Functions

Provision Authorization Policies: [VALUE ADDED]

End Point: provision {access control policy/authorization data}

Supporting Service Genre(s): authenticate (agent), authorize {access control policy/authorization data}, authorize {business rule/policy data}, activity log {workflow log}

Supporting Service Usage Model(s): identity

Primary Resource: access control policy/authorization data

Secondary Resource(s): authentication data, business rule/policy data, workflow log

Object(s): authorization policies

Provision Business Rules: [VALUE ADDED]

End Point: provision {business rule/policy data}

Supporting Service Genre(s): authenticate (agent), authorize {access control policy/authorization data}, authorize {business rule/policy data}, activity log {workflow log}

Supporting Service Usage Model(s): identity

Primary Resource: business rule/policy data

Secondary Resource(s): authentication data, access control policy/authorization data, workflow log

Object(s): business rules

Provision Authentication: [VALUE ADDED]

End Point: provision {authentication data}

Supporting Service Genre(s): authenticate (agent), authorize {access control policy/authorization data}, authorize {business rule/policy data}, activity log {workflow log}

Supporting Service Usage Model(s): identity

Primary Resource: authentication data

Secondary Resource(s): access control policy/authorization data, business rule/policy data, workflow log

Object(s): authentication data

Provision Users: [VALUE ADDED]

End Point: provision {identity data}

Supporting Service Genre(s): authenticate (agent), authorize {access control policy/authorization data}, authorize {business rule/policy data}, activity log {workflow log}

Supporting Service Usage Model(s): identity, identity provisioning and management

Primary Resource: identity data

Secondary Resource(s): authentication data, access control policy/authorization data, business rule/policy data, identifier resolution data, workflow log

Object(s): identity data

Provision System Data: [VALUE ADDED]

End Point: provision {system data}

Supporting Service Genre(s): authenticate {agent}, authorize {access control policy/authorization data}, authorize {business rule/policy data}, activity log {workflow log}

Supporting Service Usage Model(s): identity

Primary Resource: system data

Secondary Resource(s): authentication data, access control policy/authorization data, business rule/policy data, identifier resolution data, workflow log

Object(s): system provision data

2.12.2 Registry Management Functions

Adjust Registry Search Index: [BASE]

End Point: index {registry}

Supporting Service Genre(s): authenticate {agent}, authorize {access control policy/authorization data}, authorize {business rule/policy data}, index {metadata registry}, index {repository registry}, activity log {workflow log}

Supporting Service Usage Model(s): identity

Primary Resource: registry (metadata registry, repository registry), index data

Secondary Resource(s): authentication data, access control policy/authorization data, business rule/policy data, identifier resolution data, system data, workflow log

Object(s): indexing parameters

Compute Content Object Discovery Ranking: [VALUE ADDED]

End Point: generate metadata {content object ranking metadata object}

Supporting Service Genre(s): authenticate {agent}, authorize {access control policy/authorization data}, authorize {business rule/policy data}, activity log {workflow log}

Supporting Service Usage Model(s): identity

Primary Resource: metadata registry

Secondary Resource(s): authentication data, access control policy/authorization data, business rule/policy data, identifier resolution data, index data, workflow log

Object(s): content object

Compute Content Object Discovery Rating: [VALUE ADDED]

End Point: generate metadata {content object rating metadata object}

Supporting Service Genre(s): authenticate {agent}, authorize {access control policy/authorization data}, authorize {business rule/policy data}, activity log {workflow log}

Supporting Service Usage Model(s): identity

Primary Resource: metadata registry

Secondary Resource(s): authentication data, access control policy/authorization data, business rule/policy data, identifier resolution data, index data, workflow log

Object(s): content object

Compute Content Object Discovery Classifications: [VALUE ADDED]

End Point: generate metadata {content object classification metadata object}

Supporting Service Genre(s): authenticate {agent}, authorize {access control policy/authorization data}, authorize {business rule/policy data}, activity log {workflow log}

Supporting Service Usage Model(s): identity

Primary Resource: metadata registry

Secondary Resource(s): authentication data, access control policy/authorization data, business rule/policy data, identifier resolution data, index data, workflow log

Object(s): content object

Determine Metadata Object Discovery Recommendations: [VALUE ADDED]

End Point: generate metadata {content object recommendation metadata object}

Supporting Service Genre(s): authenticate {agent}, authorize {access control policy/authorization data}, authorize {business rule/policy data}, activity log {workflow log}

Supporting Service Usage Model(s): identity

Primary Resource: metadata registry

Secondary Resource(s): authentication data, access control policy/authorization data, business rule/policy data, identifier resolution data, index data, workflow log

Object(s): content object

Compute Relationships: [VALUE ADDED]

End Point: generate metadata {content object relationship metadata object}

Supporting Service Genre(s): authenticate {agent}, authorize {access control policy/authorization data}, authorize {business rule/policy data}, activity log {workflow log}

Supporting Service Usage Model(s): identity

Primary Resource: metadata registry

Secondary Resource(s): authentication data, access control policy/authorization data, business rule/policy data, identifier resolution data, index data, workflow log

Object(s): content object

Monitor Registry Behaviour: [VALUE ADDED]

End Point: monitor {registry}

Supporting Service Genre(s): authenticate {agent}, authorize {access control policy/authorization data}, authorize {business rule/policy data}, activity log {workflow log}

Supporting Service Usage Model(s): identity

Primary Resource: registry (metadata registry, repository registry, collection registry)

Secondary Resource(s): authentication data, access control policy/authorization data, business rule/policy data, identifier resolution data, index data, system data, workflow log

Inspect Registry: [VALUE ADDED]

End Point: inspect {registry}

Supporting Service Genre(s): authenticate {agent}, authorize {access control policy/authorization data}, authorize {business rule/policy data}, activity log {workflow log}

Supporting Service Usage Model(s): identity, browse

Primary Resource: registry (metadata registry, repository registry, collection registry)

Secondary Resource(s): authentication data, access control policy/authorization data, business rule/policy data, identifier resolution data, index data, system data, workflow log

Analyse Registry: [VALUE ADDED]

End Point: analyze {registry}

Supporting Service Genre(s): authenticate {agent}, authorize {access control policy/authorization data}, authorize {business rule/policy data}, activity log {workflow log}

Supporting Service Usage Model(s): identity

Primary Resource: registry (metadata registry, repository registry, collection registry), index data, workflow log, query log

Secondary Resource(s): authentication data, access control policy/authorization data, business rule/policy data, identifier resolution data, system data

Backup Registry: [BASE]

End Point: backup {registry}

Supporting Service Genre(s): authenticate {agent}, authorize {access control policy/authorization data}, authorize {business rule/policy data}, activity log {workflow log}

Supporting Service Usage Model(s): identity

Primary Resource: registry (metadata registry, repository registry, collection registry)

Secondary Resource(s): authentication data, access control policy/authorization data, business rule/policy data, identifier resolution data, index data, system data, workflow log

2.13 Design Decisions & Tradeoffs

The Design Decisions & Tradeoffs element documents overall choices, tradeoffs and their implications on the design of the Service Usage Model itself. It does not address the design of applications or the internal implementation of the elements of the Service Usage Model. No critical or essential information required to understand the Service Usage Model should be included.

Capabilities and features of the federated metadata registry provisioning and management service usage model are based on a core set of features and requirements (refer to <http://www.cordra.net>). Specific limitations are defined in the *Applicability* element of the Service Usage Model Definition.

Only essential functions needed to provide basic provisioning and management capabilities are considered core features of the federated metadata registry provisioning and management service usage model. Features and functionality that could be layered on core features or built as composite services are labelled as value added. This distinction need not be carried directly into an implementation.

The service genres used in this service usage model provide service end points used to support the provisioning and management of the central registry. Providing a service end point was the criterion for defining a service genre. Many of the service genres described map directly to one of the exposed functionalities of the repository federation. These service genres could also be expressed as small service usage models that capture all of the workflow and logic presented in the functionality and which are built on smaller specific service genres. The decision to use service genres instead of small service usage models permits building small reusable components rather than components that encode more workflow and processing.

Most of the service genres are defined to provide a single service end point. More complex service genres that provide a collection of related behaviours could be defined as an alternative approach. The decision to use multiple service genres again permits building

small components rather than components that combine potentially different behaviours. In some cases, closely related functionality has been combined in a single service genre to eliminate creating a proliferation of service genres that are almost identical.

The federated metadata registry provisioning and management service usage model assumes services are part of a trust environment and that full authentication controls are not required to use services within the trust environment. Authorization processes are still required within the trust environment, and to gain initial access to the trust environment. Providing the trust environment of components simplifies the management of authentication and authorization by individual services and for individual users but opens the potential for components to be improperly used outside of the trust environment. Such use will compromise the integrity of a federated metadata registry.

The federated metadata registry provisioning and management service usage model has not been specialised to the needs of a particular community. An individual community of practice will need to specialise the service usage model with their specific data models, policies and rules.

Specific Service Genre Modelling Decisions

None.

2.14 Implementation Guide & Dependencies

The Implementation Guide & Dependencies element describes issues of concern in implementing applications that use the Service Usage Model, including organization, performance, behaviours, representations, policies. Resolution of issues discussed is deferred to the actual application design. No critical or essential information required to understand the Service Usage Model should be included.

The federated metadata registry provisioning and management service usage model does not specify the design or details of the implementation. The implementation guidelines apply to a general federated metadata registry application that provides the functionality described.

Value-added capabilities could be built by service composition. A federated metadata registry application may directly implement more than the core features or may directly implement value-added capabilities. Direct implementation may result in improved performance.

For authorization and authentication, the federated metadata registry provisioning and management service usage model assumes a federated metadata registry application will integrate with an existing identity management system. Identity management is described in the identity service usage model and identity provisioning and management service usage model.

The federated metadata registry provisioning and management service usage model assumes a federated metadata registry application will integrate with an existing identifier system. Management of identifiers and assignment of identifiers to objects are considered part of the details of the service genres within the identifier service usage model. Provision and management of identifiers are detailed within identifier provisioning and management service usage model. Explicit functionality describing identifier management and assignment is not described herein.

Several services operate on sets or batches of data. Details of batch process (including flow control) and transaction processing are deferred to the federated metadata registry applications and the service expressions that specialise the service genres.

The federated metadata registry provisioning and management service usage model does not specify details of the infrastructure or components required to implement a federated metadata registry application. A federated metadata registry instance may require distributed or

replicated servers and multiple service instances and resource sets for load balancing and performance.

The federated metadata registry provisioning and management service usage model assumes content within the federated metadata registry application is stored and managed by enterprise-level data management components.

The federated metadata registry provisioning and management service usage model assumes content within the federated metadata registry application uses a formal indexing process to aid in discovery and search. The primary discovery process is through the formal index data.

Some steps (computing recommendations, reindexing) may be performed asynchronously with user requests. The implementation should be designed to ensure consistency of behaviour and results.

2.15 Applicable Standards

The Applicable Standards element lists domain-specific standards required to implement the Service Usage Model as a whole. Standards are described in terms of name, version and citation link. Conformance requirements and extensions should be noted. Standards used to implement applications are excluded. No critical or essential information required to understand the Service Usage Model should be included.

None. No standards are directly applicable to the federated metadata registry provisioning and management service usage model as a whole.

The service genres that are a part of the federated metadata registry provisioning and management service usage model SHALL be defined in terms of applicable standards and specifications. Applicable classes of standards and examples include:

- Repository and Registry Description Standards: Standards that define the data models used to describe a registry, repository or content collection, e.g., ISO 2146.
- Content Object Metadata Standards: Standards that describe content object metadata, e.g., DC, LOM, MODS.
- Identifier Standards: Standards used to identify registry objects and content objects, e.g., HANDLE, OpenURI, Info.
- Authentication and Authorization Standards: Standards supporting authorization and authentication of users and requests, e.g., SAML, XACML.

2.16 Known Uses

The Known Uses element documents actual uses of the Service Usage Model in applications and systems, including how used, extensions, limits.

Actual: None as documented.

Potential Service Usage Model:

A service-based implementation of the ADL-Registry (ADL-R) would be an example of a registry and repository federation that follows this service usage model.

Potential Application: The TILIS content registry is an example of a registry and repository federation that follows this service usage model.

Potential Application: The ADL-Registry (ADL-R) is an example of a registry and repository federation that follows this service usage model.

2.17 Resources Used

The Resources element documents all Resources used in the Service Usage Model. Resources are identified by name and version.

The service genres and service usage models within the repository federation service usage model manage the resources. The resources needed to provide essential functionality are labelled [BASE] while others are labelled [VALUE ADDED]. Identified resources include:

metadata registry: [BASE]

Managed storage for metadata objects describing the content objects exposed through the repository federation.

collection registry: [BASE]

Managed storage for collection objects describing the collections exposed through the repository federation.

repository registry: [BASE]

Managed storage for repository objects describing the repositories exposed through the repository federation.

system data: [BASE]

Managed storage for system-level objects used to describe, manage and control the registry (vocabularies, schemata, registry descriptions).

access control policy/authorization data: [BASE]

Managed storage for access control policies used to control access to objects managed by the registry.

business rule/policy data: [BASE]

Managed storage for policy and business rules used to control operations of the registry.

index data: [BASE]

Indexing information for the metadata registry used in search and discovery.

identifier resolution data: [BASE]

Identifier attributes used in identifier resolution and object access.

identity data: [BASE]

Information about user identities for those who access the registry.

authentication data: [BASE]

Information about authentication of users for those who access the registry.

workflow log: [VALUE ADDED]

Logs of registry operations.

query log: [VALUE ADDED]

Logs of registry queries and search results.

2.18 Services (Service Genres) Used

The Service Genres Used element documents all Service Genres used in the Service Usage Model. Service Genres are identified by name and version.

The federated metadata registry provisioning and management service usage model includes the service genres listed below. The service genres needed to provide essential functionality are labelled [BASE] while others are labelled [VALUE ADDED]. The service genres that are required for internal operations within the repository federation service usage model but not exposed to end users as service end points are labelled [INTERNAL].

NB: These internal service genres MAY provide useful end user functionality for other applications, but such uses are not considered as part of the federated metadata registry provisioning and management service usage model.

provision: Vx.xx.

[\[link to service genre\]](#)

Add and manage data and settings to control overall registry behaviour.

- provision {system data}: [BASE]
Add and manage system control data and settings in the system data resource to control overall registry behaviour.
- provision {access control policy/authorization data}: [BASE]
Add and manage access control policies in the access control policy/authorization data resource. This service genre is part of the authorization service usage model.
- provision {business rule/policy data}: [BASE]
Add and manage business rules and policies in the business rule/policy data resource.
- provision {identity data}: [BASE]
Add and manage identity data in the identity data resource. This service genre is part of the identity service usage model.
- provision {authentication data}: [BASE]
Add and manage authentication data in the authentication data resource. This service genre is part of the identity service usage model.

index: Vx.xx.

[\[link to service genre\]](#)

Index or reindex a registry to support discovery from the registry.

- index {metadata registry}: [VALUE ADDED]
Index or reindex the metadata registry to support discovery from the metadata registry. Internally, the index service genre manages the index data resource.
- index {collection registry}: [VALUE ADDED]
Index or reindex the collection registry to support discovery from the collection registry. Internally, the index service genre manages the index data resource.
- index {repository registry}: [VALUE ADDED]
Index or reindex the repository registry to support discovery from the repository registry. Internally, the index service genre manages the index data resource.
- index {registry}: [VALUE ADDED]
Index/reindex all the core registry components. Internally, the index service genre manages the index data resource.

inspect: Vx.xx.

[\[link to service genre\]](#)

Inspect a registry. Query for registry settings and values.

- inspect {registry}: [VALUE ADDED]
Inspect the registry.

monitor: Vx.xx.

[\[link to service genre\]](#)

Monitor the behaviour and performance of a registry. Query for performance and behaviour data results.

- monitor {registry}: [VALUE ADDED]
Monitor the registry.

backup: Vx.xx.

[\[link to service genre\]](#)

Backup a registry.

- backup {registry}: [VALUE ADDED]
Backup the core registry components.

analyze: Vx.xx.

[\[link to service genre\]](#)

Analyse a registry.

- analyze {registry}: [VALUE ADDED]
Obtain analytical data from a registry and prepare registry analytics.

authenticate: Vx.xx.

[\[link to service genre\]](#)

Authenticate a user (including a service acting on behalf of a user) using information from the authentication data resource.

- authenticate {agent}: [BASE/INTERNAL]
Authenticate a user agent (or a service acting on behalf of a user agent) using information from the authentication data resource. This service genre is part of the identity service usage model.

authorize: Vx.xx.

[\[link to service genre\]](#)

Authorise a user (or a service on behalf of a user) to perform a particular behaviour or access a particular resource using business rules and policies.

- authorize {access control policy/authorization data}: [BASE/INTERNAL]
Authorise a user (or a service acting on behalf of a user) to perform a particular behaviour or access a particular resource using access control policies in the access control policy/authorization resource. This service genre is part of the authorization service usage model.
- authorize {business rule/policy data}: [BASE/INTERNAL]
Authorise a user (or a service on behalf of a user) to perform a particular behaviour or access a particular resource using business rules and policies in the business rule/policy resource.

generate metadata: Vx.xx.

[\[link to service genre\]](#)

Generate an object based on information available in another object.

- generate metadata {content object classification metadata object}: [VALUE ADDED]
Generate/compute classification metadata for a content object. This service genre is applied to a content object.
- generate metadata {content object ranking metadata object}: [VALUE ADDED]
Generate/compute ranking metadata for a content object. This service genre is applied to a content object.
- generate metadata {content object rating metadata object}: [VALUE ADDED]
Generate/compute rating metadata for a content object. This service genre is applied to a content object.
- generate metadata {content object recommendation metadata object}: [VALUE ADDED]
Generate/compute recommendation metadata for a content object. This service genre is applied to a content object.
- generate metadata {content object relationship metadata object}: [VALUE ADDED]
Generate/compute relationship metadata for content objects. This service genre is applied to a content object.

schema transform: Vx.xx.

[\[link to service genre\]](#)

Apply a schema transform to an object.

- schema transform {metadata object}: [VALUE ADDED/INTERNAL]
Transform a metadata object from one representation to another.
- schema transform {collection object}: [VALUE ADDED/INTERNAL]
Transform a collection object from one representation to another.
- schema transform {repository object}: [VALUE ADDED/INTERNAL]
Transform a repository object from one representation to another.
- schema transform {content object}: [VALUE ADDED/INTERNAL]
Transform a content object from one representation to another.

render transform: Vx.xx.

[\[link to service genre\]](#)

Apply a render transform to an object.

- render transform {metadata object}: [BASE/INTERNAL]
Apply a rendering transform (XSLT) to a metadata object.
- render transform {collection object}: [BASE/INTERNAL]
Apply a rendering transform (XSLT) to a collection object.
- render transform {repository object}: [BASE/INTERNAL]
Apply a rendering transform (XSLT) to a repository object.
- render transform {content object}: [BASE/INTERNAL]
Apply a rendering transform (XSLT) to a content object.

filter: Vx.xx.

[\[link to service genre\]](#)

Filter (via provided data) the object (or set of objects) as part of an access control workflow.

- filter {metadata registry}: [VALUE ADDED/INTERNAL]
Filter (via authorization data) the metadata object (or a set of objects) retrieved from the metadata registry as part of a metadata object access workflow. Internally, the filter service genre uses the access control policy/authorization data resource.
- filter {collection registry}: [VALUE ADDED/INTERNAL]
Filter (via authorization data) the collection object (or a set of objects) retrieved from the collection registry as part of a collection object access workflow. Internally, the filter service genre uses the access control policy/authorization data resource.
- filter {repository registry}: [VALUE ADDED/INTERNAL]
Filter (via authorization data) the repository object (or a set of objects) retrieved from the repository registry as part of a repository object access workflow. Internally, the filter service genre uses the access control policy/authorization data resource.
- filter {repository}: [VALUE ADDED/INTERNAL]
Filter (via authorization data) the repository object (or a set of objects) retrieved from a repository as part of a repository access workflow. Internally, the filter service genre uses the access control policy/authorization data resource.

log: Vx.xx.

[link to service genre]

Log an activity or action in a log object.

activity log {workflow log}: [VALUE ADDED/INTERNAL]

Log an activity/action in the workflow log.

2.19 Related Service Usage Models

The Related Service Usage Models element documents and illustrates how the Service Usage Model is related to other Service Usage Models required to provide its stated behaviours. Related Service Usage Models are identified by name and version. No critical or essential information required to understand the Service Usage Model should be included.

The service usage models listed below are assumed to exist.

identifier: [BASE] Vx.xx.

[link to service usage model]

The identifier service usage model is an integrated set of service genres that provides the operational identifier infrastructure used by the federated metadata registry provisioning and management service usage model. Functionality includes the creation, registering, publishing, managing and resolving of individual identifiers.

identifier provisioning and management: [BASE] Vx.xx.

[link to service usage model]

The identifier provisioning and management service usage model is an integrated set of service genres that manages the identifier infrastructure used by the federated metadata registry provisioning and management service usage model. Functionality includes provisioning and management of the identifier system.

repository storage manager: [BASE] Vx.xx.

[link to service usage model]

The storage manager service usage model is an integrated set of service genres that provides access operators (CRUD, versioning, audit, ACID transactions) for a stored data collection, e.g., the metadata registry, ~~collection registry~~, repository registry, ~~archival repository~~, ~~eserow repository~~.

identity: [BASE] Vx.xx.

[\[link to service usage model\]](#)

The identity service usage model is an integrated set of service genres that provides the operational user identity infrastructure used by the federated metadata registry provisioning and management service usage model. Functionality includes the creation of user right and roles, user authentication and user authorization.

NB: Authentication and authorization are combined within a single service usage model.

identity provisioning and management: [BASE] Vx.xx.

[\[link to service usage model\]](#)

The identity provisioning and management service usage model is an integrated set of service genres that manages the user identity infrastructure used by the federated metadata registry provisioning and management service usage model. Functionality includes provisioning and management of the identity system.

NB: Authentication and authorization are combined within a single service usage model.

browse: [VALUE ADDED] Vx.xx.

[\[link to service usage model\]](#)

The browse service usage model provides an integrated set of service genres used to develop browser-based interfaces for a repository federation.

repository federation (content): V1.0.0

[\[link to service usage model\]](#)

The repository federation (content) service usage model provides an integrated set of service genres used to populate and use the metadata registry that supports a repository federation. Functionality includes content management (creation and management of metadata objects within a repository federation), content discovery (discovery of content objects from a repository federation) and content delivery (retrieval of and access to content objects discovered through a repository federation).

2.20 Service Patterns/CORE Service Usage Models Used

The Service Patterns/CORE Service Usage Models element documents all CORE Service Usage Models actually used in the Service Usage Model. CORE Service Usage Models are identified by name and version.

Actual: None as documented.

Potential: AVTEROFTL workflow, Vx.xx.

[\[link to CORE service usage model\]](#)

Many of the registry operations follow a common workflow pattern that could be abstracted. The AVTEROFTL workflow core service usage model defines this pattern.

Steps in this commonly reoccurring pattern include:

- Authenticate and authorise: authenticate and apply business rules to validate the request/operation [A].

- Validate: validate the data object associated with the request/operation [V].
- Transform: apply a (schema) transform to the data object associated with the request [T].
- Enhance: enhance or modify the data object associated with the request (adding data not in the original request) [E].
- Route: route the request to a specific operator [R].
- Operate: perform the requested operation and obtain results [O].
- Filter: filter the results (apply access controls and other filters to the results) [F].
- Transform: apply a (schema) transform to the results data object [T].
- Render: apply a rendering transform to the results data object [R].
- Log: log the activity [L].

Of these steps only AVROL (Authenticate and Authorise; Validate; Route; Operate; Log) are compulsory and recur for every instance of the workflow. The remaining steps are all optional, and are included in the workflow only as required by the business logic. If steps are omitted, then the data input to the workflow is what is required as input to the operation, and the operation output produces what is required by the workflow consumer, with no further modification.

Potential: ROAP obtain, Vx.xx.

[\[link to CORE service usage model\]](#)

The ROAP obtain core service usage model is a combination of four services used to access an object given an object identifier. Steps in this commonly reoccurring pattern include:

- A request, in the form of an identifier, is resolved (including multiple resolution based on a FRBR model to obtain the appropriate work, expression, manifestation or instance) yielding the source resource that manages the object [R].
- An obtain accessor is used to get the object [O].
- Subject the object to access control or obligation filtering [A].
- The resulting object dissemination is encoded and packaged in the requested format [P].

Search is used as a precursor to the ROAP obtain process to discover the object identifier. Selected object identifiers are then passed through the ROAP obtain workflow to obtain the object. The ROAP obtain service usage model itself is an operator in an AVTEROFTL workflow.

2.21 References

The References element includes references and bibliographic citations to works needed to understand the Service Usage Model.

None

Glossary

The Glossary element defines domain-specific terms used in documenting the Service Usage Model.

collection: A named, managed set of content objects. A collection may span multiple repositories. A collection is a type of resource.

collection object: Descriptive metadata, objective or subjective, for a collection. A collection object is an object managed by a collection resource.

content object: An object, digital or physical; part of a managed collection of content objects.

FRBR: Functional Requirements for Bibliographic Records: A standard for the modelling of entities in library cataloguing, and for the relations between them. The FRBR ontology of entities is used in the Service Usage Model to differentiate content objects which may be regarded as the same at some underlying levels but not others.

FRBR Work: An abstract representation of a content object as a presentation of content. The most general representation of a element of content, independent of its presentation, encoding or storage location. Two content objects present the same FRBR Work if their content is deemed to be underlying the same, ignoring revisions and reformulations. Two content objects belong to distinct FRBR Works if their content is significantly different.

FRBR Expression: An abstract representation of a content object as a presentation of content. Two content objects present the same FRBR Expression if their content is identical, and the objects differ only in the presentation of that content (including the file format). Two content objects are distinct FRBR expressions of the same FRBR work if the content of one is a minor modification of the content of the other (e.g., different drafts or versions).

FRBR Manifestation: A representation of a content object as a presentation of content. Two content objects present the same FRBR Manifestation if their content is identical and the presentation and storage and representational encoding of the content is identical—i.e., they are they same file, though the location and access metadata of the objects may be different. Two content objects are distinct FRBR Manifestations of the same FRBR expression if their content is identical, but the presentation of the content differs (e.g., different file encoding, different file format).

FRBR Instance: A content object viewed in the context of the FRBR model, bound to a specific single storage location. All content objects are FRBR instances; all content objects can be mapped to a specific FRBR Manifestation, FRBR Expression, and FRBR Work, as increasingly abstract representations of the object content.

repository federation: A managed set of repositories, containing content objects (that are within collections) and a registry used to store metadata (metadata objects) used for discovery.

metadata object: Descriptive metadata, objective or subjective, for a content object.

registry: A repository containing metadata (metadata objects) for content objects from a set of repositories. In the federation model described, the registry also contains collection objects and repository objects. A registry is a type of resource.

metadata registry: A repository containing metadata about content objects, i.e., metadata objects.

collection registry: A repository containing metadata about collections, i.e., collection objects.

repository registry: A repository containing metadata about repositories, i.e., repository objects.

registry object: A metadata object describing a registry. A registry object is an object managed by a registry resource.

repository: A storage and management system for one or more sets of content objects or metadata objects. A repository is a type of resource.

repository object: A metadata object describing a repository. A repository object is an object managed by a repository resource.

user: Any application or client accessing a service. The client or application may be a proxy for a human user / end user.