LMS-ePortfolio Interoperability Project

Requirements and Architecture Consultation Document

Document Information	
Project Name	LMS-ePortfolio Interoperability Project
Document Title	Requirements and Architecture Consultation Document
Author	Sandy Britain
Date	30/10/10
Version	1
Comments	Includes revisions based on the consultation meeting 8/11/10.
Contact	sandybritain@gmail.com

Table of Contents

Executive Summary

Intended readership for this Document

This document is for both non-technical and technical stakeholders in the Schools MLE Programme with an interest in defining an approach to technical interoperability between the Mahara ePortfolio system (implemented as a national service under the name of MyPortfolio) and school systems (particularly LMS systems).

The Purpose of this Document

The purpose of the document is to outline the context of the project, stakeholder concerns and a proposal for a technical architecture to support interoperability within the scope of the project.

The Contents of the Document

Section 1 of the document contains an introduction to the rationale for interoperability between LMS systems and ePortfolios. It also scopes the interoperability requirements.

Section 2 presents a summary of stakeholder concerns.

Section 3 presents some scenarios of use which serve to make explicit the patterns of use by students and teachers

Section 4 draws out a set of functional requirements for interoperability and suggests some general architectural principles to guide development

Section 5 Looks in more detail at a technical approach and is mainly intended for developers and those with an interest in the technical details. For non-technical readers it contains an overview of the design issues under consideration.

Section 6 identifies the next steps and sets out some recommendations for further work

1. Project Background and Scope

In recent years ePortfolios have begun to emerge as an important component of the technology landscape within education institutions at both school and tertiary levels. Whilst the purpose of an ePortfolio can vary considerably depending on the educational context and level, there is increasing consensus that the ePortfolio fulfills a need for a student-owned and student-centred space to manage, record and present their work, reflections and achievements that LMS systems cannot easily cater for.

Within the Schools MLE Programme the LMS system is seen as a 'hub' component in the curriculum and pedagogy space. ePortfolios represent one of the key satellite systems in this space. The Ministry of Education (MOE) conducted an initial investigation into ePortfolio use in schools in 2009.[1] One of the primary questions motivating that work was how best to support schools in developing their use of ePortfolios. The report recommended that there was little benefit in supporting a proliferation of ePortfolio systems, the main issue was to support one that had the potential to interoperate well as part of an MLE with other school systems.

The open source ePortfolio platform originally developed in New Zealand presented itself as an obvious choice, not least because it offers an implementation of the preferred candidate for an ePortfolio interoperability standard (LEAP 2A). Since then a decision has been taken to support certain aspects of the ongoing development of Mahara as a national reference system for ePortfolios and to provide an implementation of Mahara (currently under the name of MyPortfolio) as a hosted service to schools. Under this scheme every school student in New Zealand is eligible for a MyPortfolio account which they can use throughout their school career.

With these developments as a background, the main question, in the context of the MLE programme, now becomes: how can MyPortfolio and school LMS systems be best integrated with each other to operate effectively and cohesively together? This is a multi-faceted question which can be partly solved through technical interoperability but also includes social and organisational dimensions around the rules and patterns of use of these systems. The ePortfolio-LMS integration project will explore user requirements for integration of Mahara with the LMS systems and work with developers to establish an architectural model for solving the technical interoperability requirements identified.

1.2 Project Context and System Components

The diagram (figure 1) below illustrates the main system components under consideration here and the current interactions both between system components and with human users. The figure represents a general schematic of the current situation typically found in schools using both an LMS system and MyPortfolio. The figure draws a distinction between those systems that are 'owned' and managed by the school (those above the dashed dividing line) and those that exist outside of institutional ownership and control.

1 This work resulted in the report: 'ePortfolios: Celebrating Learning' (Fox, Britain and Hall, 2009).

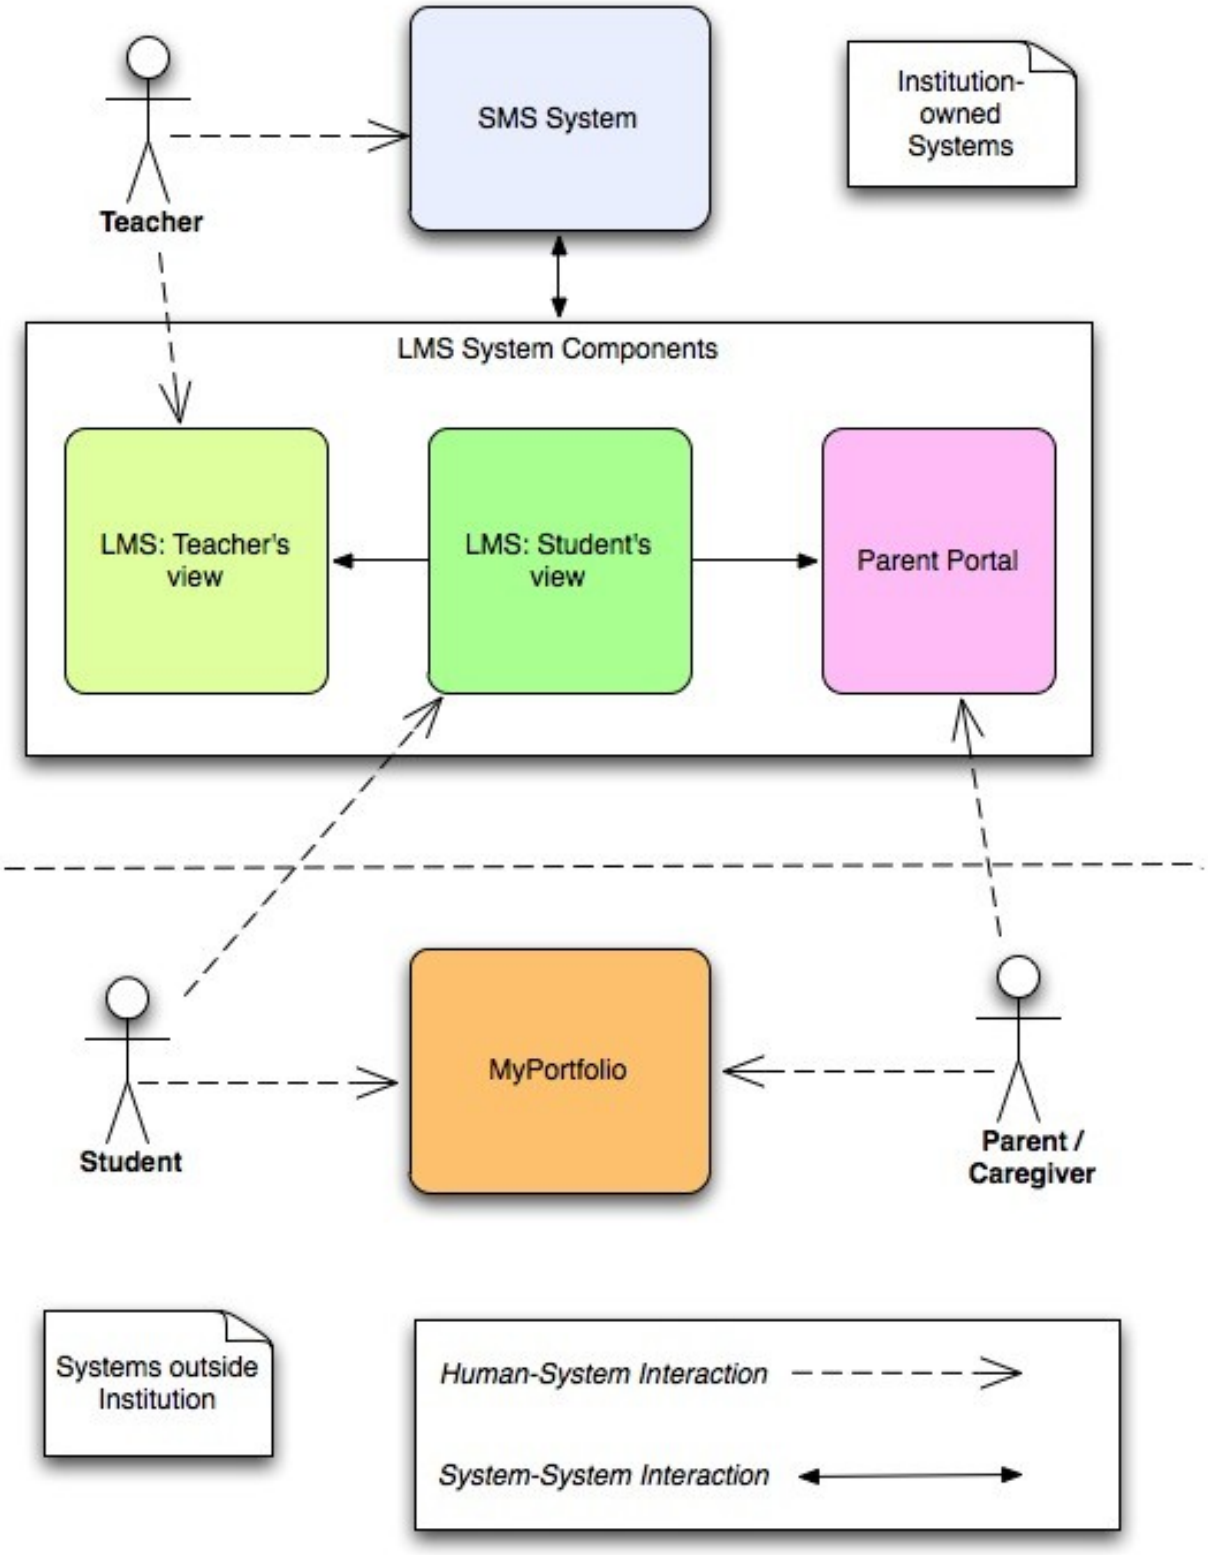

Figure 1: Overview of System Components and Current Interactions

The system components and interactions are described below. Within the school-owned space the primary system of concern here is the LMS system. There are many possible ways to describe the components of an LMS, but here we depict the various role-based views that are available. Naturally the teacher has a view onto the system with various teacher-oriented tools and spaces (e.g. activity creation and management tools, course management tools, assessment tools etc). The students have a different view which is typically limited to their engagement with the courses and activities set-up by the teacher and interactions with their peers in class-based discussions and so on. Within New Zealand schools an important third component or view is the Parent Portal. In some cases this forms part of the LMS itself (e.g with Ultranet and KnowledgeNet) as depicted here, or it may be an independent piece of software. The Parent Portal provides a view onto the students work and progress for the parents or caregivers to view and integrates NZQA assessment information from the SMS system and items from the LMS system.

MyPortfolio (as an instance of the Mahara platform which is provided as an independently-hosted national service) sits outside of the school systems. The key idea is that MyPortfolio is a student-owned and managed space which allows students to organise their work according to their own categories and projects, to engage in reflective study and to create selective views for others such as parents, teachers, peers, potential employers and so on to see and/or provide feedback.

In some cases parents can currently access a view of their child's MyPortfolio space through a 'secret-url' mechanism, but in others parents are encouraged only to access information through the Parent Portal. Although not shown on the diagram for simplicity, teachers can also be provided with access to MyPortfolio views as teachers will have their own accounts on the system, which means that students can provide peers and teachers with views to their portfolio without resorting to the 'secret-url' mechanism.

What is missing from this current picture is any system-to-system interaction between the school-owned systems (particularly LMS systems) and MyPortfolio. This current lack of interoperability has been identified as a hindrance to the effective use of MyPortfolio in a schools context and provides the focus for the current work.

1.3 Goals and Drivers of the project:

The key driver for the project is to increase the value of MyPortfolio as hosted service to NZ schools by developing an interoperability solution with existing school LMS systems which meets the needs of stakeholders. The high-level goals of the project are:

- To identify and document stakeholder concerns from both schools and relevant software developers around interoperability between LMS systems and MyPortfolio
- To identify functional requirements for interoperability
- To investigate the suitability of the LEAP2A specification for data transport between systems
- To create an architectural model which meets concerns and handles requirements
- To achieve consensus amongst LMS and mahara developers on the architecture and technical implementation plan

1.4 Scope

The scope of the project is confined to modeling interoperability between LMS systems (and other related school MLE components such SMS and Parent Portals) and MyPortfolio. It does not extend to interoperability between different ePortfolio systems, (that problem is being worked on elsewhere – see the JISC PIOP projects[2]), although the developments arising from this project may also support more general interoperability between Mahara and other systems.

There are five primary aspects of the technical interoperability under consideration. These are:
1. Content Transfer. This is the most obvious driver for interoperability – i.e. the ability for a student to transfer some digital artifact(s) or personal information from their LMS system into their ePortfolio and vice versa.
2. Single Sign On. SSO (in some form) is a key aspect of interoperability both for increasing usability for the learner – to prevent them having separate log-ins, but also for secure data transfer between systems it is necessary for the remote system to be able to establish the credentials of the person initiating the transfer. The same mechanism may or may not be used to support both cases.
3. User Provisioning. Some level of automated provisioning of user accounts into MyPortfolio is important to prevent the need for manual account creation. This is closely tied to SSO.
4. Notfications. It is important, where the content of MyPortfolio may be accessed through an external participating system (e.g. a Parent Portal) to provide notifications of updates to content or views. Other forms of notification may also be required.
5. Theming. For some users it is important for the ePortfolio to be branded or themed in the same way as the other school systems. For others it is important to be able to distinguish the two. Whilst this is a consideration here, it sits apart from the other main aspects of interoperability and is of lower priority.[3]

The diagram in figure 2 below shows the same systems as figure 1 with the scoped system-to-system interactions between MyPortfolio and school-owned systems included. Note that theming is not included here as it is not being considered as an automated system-to-system interaction. There are some other key points to note about this diagram:

- The Single Sign On (SSO) mechanism is depicted as uni-directional from the LMS environment to MyPortfolio. This already exists for schools using Moodle and MyPortfolio and for schools using the SAML-based IAM mechanism created for the MLE programme.
- Content Transfer between the students' space in the LMS and their MyPortfolio account is depicted as bi-directional. There needs to be a system-to-system authentication and authorisation method to support this which will differ from the human-to-system SSO mechanism
- Account Provisioning in MyPortfolio is uni-directional from the LMS (as a school-owned

2 The JISC ePortfolio Interoperability Programme (PIOP) http://www.jisc.ac.uk/whatwedo/programmes/elearningcapital/piop.aspx
3 Catalyst IT / Flexible Learning who manage the MyPortfolio hosting in New Zealand can provide theming services for MyPortfolio.

system) to MyPortfolio

- Notifications of updates to views in MyPortfolio may be be provided to the teacher's space in the LMS and to parents through the Parent Portal. This channel may also provide a mechanism for transfer of the content itself thus blending informal ePortfolio content (reflections etc) about an activity with formal institutional information about the activity in the LMS. This channel is depicted as uni-directional from MyPortfolio out, although some stakeholders have expressed the concern that notifications may also need to flow the other way.

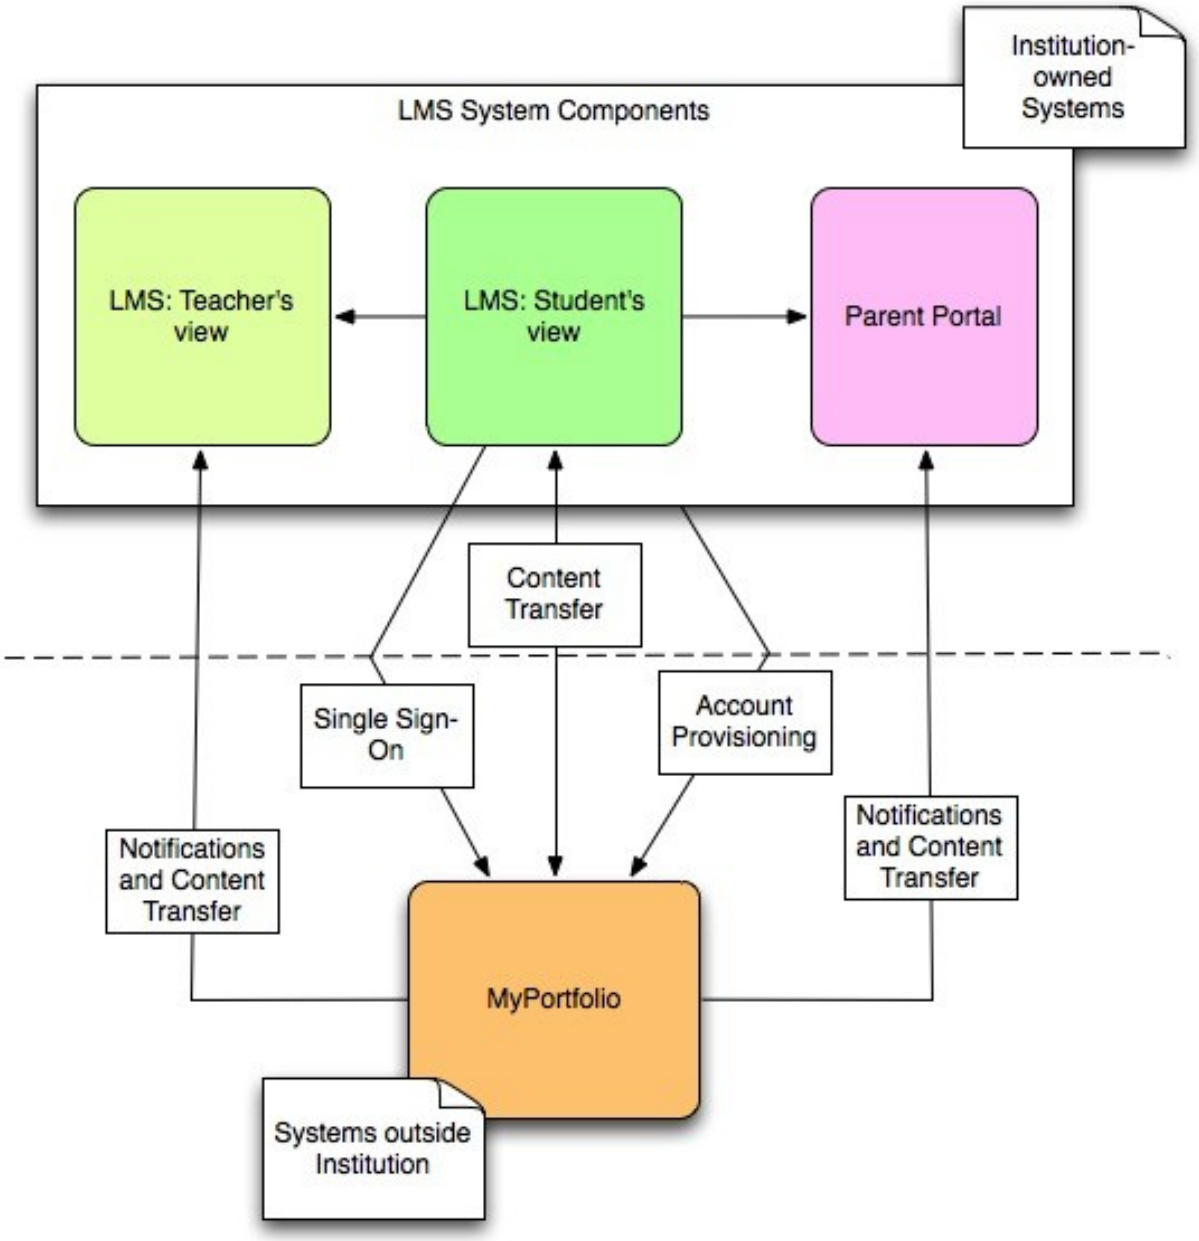

Figure 2: Scoped System Interactions

2. Stakeholder Concerns

In this section we outline a number of the concerns expressed by stakeholders during the requirements gathering phase of this work.[4] These can be broadly classified into two types:

1) Those concerns that can be directly addressed through technical interoperability in some form or which have a direct influence on the way that the technical architecture is designed or implemented
2) Those concerns which relate to social and / or organisational patterns of use of the software and do not have a direct bearing on the technical architecture.

The first category is important as those concerns are likely to lead to functional requirements for the technical work to follow. The second category is equally important as the eventual technical architecture needs not to impede the variations in practice these concerns articulate. Many of the concerns contain both types of element, hence both types of concern have been captured and an effort is made to distinguish the technical elements from the pattern of use elements.

2.1 Ownership

The question of ownership emerged as the key differentiator between ePortfolios and other school-based learning technology tools. In an ePortfolio, the student owns the space, they decide what goes into it, how it is structured and who to share it with. By contrast the LMS and ePortfolio-like activities within the context of the LMS or other school systems are 'owned' and managed by the school. Therefore there is (for example) a profound contextual difference between a reflective blogging task within the LMS and the same task within your own ePortfolio environment with respect to ownership (and responsibility for) both the artefact and the process. This concern was highlighted by almost all stakeholders interviewed and if it is treated as a top-level consideration, then it holds significant implications for the way in which other concerns may be handled from an architecture perspective: both simplifying and constraining many of the functional requirements of interoperability. These are outlined throughout the rest of the document. The fact that MyPortfolio is being provided as a national hosted service outside of any particular institution's field of responsibility serves to support the conceptual distinction being made here of student-owned and school-owned systems.

However it is important to note that in practical terms (especially where younger children are concerned) 'student-owned' does not mean there is no teacher involvement in directing ePortfolio activities or managing ePortfolio content. There is a continuum of 'student-ownership' which has to do with both the maturity of the students and the purpose of the ePortfolio in a particular educational context.

Implications of Student Ownership	
Pattern of Use	*Technical*
Students of sufficient maturity are entirely	Teachers and school IT and admin staff do not

4 A list of the stakeholders consulted is provided in Appendix 1.

responsible for the content and organisation of their ePortfolio as well as account management	have 'administrator rights' which allow them access to the students account in MyPortfolio
Teacher access is necessary for younger students in order to help them with Portfolio activities	Currently teachers can create a class portfolio using the groups function. Alternatively can use the students login for access.
Students need to be able to transfer content between their LMS environment and their ePortfolio	System-to-system content transfer using a standard exchange format (LEAP2A) and an appropriate authentication and authorisation mechanism

2.2 Assessment

This concern relates to how assessed content should be treated in the ePortfolio environment. The first point to consider here is that Mahara / MyPortfolio is not really designed to be an environment which handles formally assessed content, it is more intended to be an environment in which the student collects artifacts which are evidence of a personal learning journey for reflection upon by themselves and given over to view for feedback by others. This perspective is in keeping with the ownership distinction mentioned above. Institution-managed systems are designed to handle and verify assessment information. The LMS typically provides many tools for this purpose and the SMS is the authoritative source of Grade information. These systems are controlled by the institution. The ePortfolio is owned and managed by the student, so should not be considered an authoritative source for assessment information. What if the student were to edit the work after it had been assessed in the ePortfolio? What if they were to change the grade? Or if they were to selectively edit out less positive comments after the fact? If we accept the ownership distinction made earlier, then we might respond – “Well, so what if they do?” It is nothing that they couldn't do in other ways if they so desired and if it is not relied upon as an authoritative source on assessment then it shouldn't matter.

However it is likely that students will want to export some completed and assessed work from the LMS into their ePortfolio as evidence for presentation to third parties. One option to achieve this would be to transfer a 'locked' artefact from the LMS which can be appended and annotated with reflections and commentary but cannot actually be changed. An argument against this approach is that if the Portfolio is owned by the student, then it shouldn't really contain 'locked' artifacts that the can't edit, also it is part of the learning process to continue to improve a piece of work. Another approach that has been suggested would be to use digital signatures such as xml-signatures to ensure that a third-party can verify that signed content hasn't be subsequently amended.

Another reason to want to transfer a locked “snapshot” of artifacts at a particular point in time would be if the portfolio itself (or a view rather) were to presented for assessment. Rather than attempting to 'lock' assessable views in MyPortfolio which would again contravene the principle that the ePortfolio is student -owned and managed, one way to handle it would be to export a time-stamped, digitally signed snapshot of the view into the teacher's space in the LMS for assessment. The responsibility for maintaining verified copies of assessed work then passes (appropriately) to the institution and mechanisms to support this are out of scope of the current

work.
It has been mentioned that the use of embedded video/audio stored on third party servers (e.g YouTube) presents issues for maintaining a verified and complete copy of a students view since the links may disappear. While it is worth noting this again is a pattern of use issue which is beyond the scope of the current work.

Implications of Assessment Concerns	
Pattern of Use	*Technical*
Information in the Portfolio is student-owned and therefore can be adapted at will by the student. Therefore this information cannot be relied upon to provide a verified authoritative source of the state of a piece of work at a particular time if this is required (e.g. for assessment)	Mechanisms for copying ePortfolio views for assessment purposes to school-owned systems (LMS) are required. XML-signatures could be used to verify content hasn't changed after a certain date. (see assessment concern below)

2.3 Updating and Versioning

This concern bears some relation to the previous one but in a more general sense. Should exported content be versioned to provide a history or overwritten when an updated version is exported? The view was expressed that versioning is desirable because the existence of several versions of the same work as it progresses provides rich evidence of the learning journey. Ideally both options would be available.

Implications of Versioning	
Pattern of Use	*Technical*
System users will have to manage versions of artifacts and views which increases complexity.	The importing systems will need to be able to detect what constitutes a new version of an existing artefact or view.

2.4 Exporting and Displaying Structured Content / Activities

Some content types are simple unitary entities such as a blog post or a single file. Other pieces of content may be structured into several hyper-linked html pages, perhaps with embedded images or other linked files which also need to be transported, and with comments or other annotations. In such cases the content needs to be exported in such a way as to take all necessary additional files with it and re-link them appropriately as well as making sure all other links remain valid. In addition the remote system needs to know (or make some sensible decisions about) how to display the content in a way that makes sense.

Implications of Structured Content	
Pattern of Use	*Technical*
N/A	The export format used should be able to

	handle structured content. LEAP2A is able to do this through the use of relationship elements.

2.5 Contributions of Others and Licensing

If an LMS activity involves the collaborative input of several students, or in the case of a discussion thread, the views of several participants are expressed, then should a student be permitted to export that activity in its entirety to their ePortfolio? Should they be able to do it at all? Should they only be able to do it with the permission (either implicit or explicit) of the other participants? Currently when Moodle discussion forums are exported to Mahara, all posts other than those by the exporting user are removed. This is one way to handle the issue but all of the context of the discussion is lost. One suggestion was to apply a creative commons licensing scheme attached to certain LMS collaborative activities such as discussions that provides consent for the materials to be exported to ePortfolio environments by any of the participants.

Implications of Contributions of Others	
Pattern of Use	*Technical*
Participants in collaborative activities and discussions within the LMS may need to either implicitly or explicitly give consent for those items to be exported as artifacts to MyPortfolio	The LMS could implement an administrative switch which changes an 'exportable' property on an activity or a class of activities.

2.6 Privacy and Access

If the ePortfolio is student owned and managed, then should student have the ultimate right to grant or deny access to staff members (teachers and school IT staff) to their account information. Naturally this doesn't make sense for younger children whose ePortfolio activities are to a large extent directed and managed by the teacher and / or parents. At what point would control in managing access be handed over to the students themselves? This is a valid concern but is a pattern of use question that is out of scope of this work. No technical solution is being considered.

2.7 Notifications

In situations where parents are encouraged to use the Parent Portal as the primary point of access to their child's' work and education information, then MyPortfolio should provide a notification to the Parent Portal when a view pertaining to the parent has been created or updated. Similarly teachers expressed the need for MyPortfolio to provide Notifications through to their space in the LMS when ePortfolio views have changed in some way. Currently Mahara provides email notifications which don't contain any information about what has changed. At a minimum it need to supply the student's name, the item / view that has changed and a URI link back to the view.

The view was expressed that the notifications mechanism could easily be set-up using an Atom

feed from MyPortfolio which is consumed by a lightweight feed-reader built into the LMS / Parent Portal systems. The notification payload should be as lightweight as possible to avoid performance/scalability issues (with a single national service there could be a very high quantity of notification traffic) and secondly to reduce the need for security mechanisms to an insignificant level. This means that no actual content should be exchanged as part of the notification mechanism.

One problem with the notification scheme for parents is that unless parents have a log-in to MyPortfolio which they rarely do (for reasons of server loads and the undesirability of having to provide a service which caters for all parents nationally as well as students), they can't currently receive notifications. Normal parent access to a view is through the 'secret-url' mechanism. It was noted that the secret-url model is very weak from a security perspective. Yet allowing for parent accounts is undesirable from a management perspective. If parent access is only permitted through a school parent portal system, then an alternative scheme using delegated server-to-server authentication and authorisation (e.g.OAuth) could allow parents to receive notifications and access views in a secure environment using the parent portal as an intermediary.

An additional comment on notifications was that a one-way flow from MyPortfolio out may not be sufficient, and that perhaps MyPortfolio should accept notifications that feedback on a view has been placed, so that the student can in turn react to the update. This may be useful but is not being considered in the first phase of this work.

Implications of Notifications	
Pattern of Use	*Technical*
Parents get used to using the Parent Portal to access child's work	Notifications alert teachers and parents through the LMS / Parent Portal of updates to views.
	Notifications can use an Atom feed-based approach and should be very lightweight
Clicking on the URI creates a LEAP2A GET request which exports the view contents to the Parent Portal	Notifications should include name, view, and a URI to the view
	Parents should access views through the parent portal rather than directly through secret-url if possible. Parent Logins to MyPortfolio are undesirable from a Management perspective

2.8 Relationship with internal LMS E-Portfolio Tools

LMS systems typically provide their own ePortfolio-type tools or activities to one extent or another (e.g. KnowledgeNet, Moodle and Ultranet all do). How would MyPortfolio work with those tools? Once again the question of ownership provides a useful top-level distinction here. Where ePortfolio tools and activities are provided within the LMS, they are still part of the school-owned and managed set of tools. MyPortfolio sits outside of that. Which activities take

place in the LMS ePortfolio tools and which in MyPortfolio (if both are available) is a pattern of use question that is up to the direction of the individual teacher or institution policy.

2.9 Supporting Lifelong (Life-wide) Learning

The provision of MyPortfolio as a hosted national service provides an extremely practical foundation for the broader role of ePortfolios in supporting lifelong learning as students move between educational institutions. However this raises the question of identifying students as they move between institutions and how many accounts each student can have in MyPortfolio. Whilst this broader question is out of scope of the project, it does bear on the issue of implementing user provisioning and student account creation in each institution. If students were to be provided automatically with a new account in each new institution they move, then it would be important for them to be able to export their portfolio materials from a previous account into their new one (which can be done in Mahara using the existing LEAP2A import and export functionality) or simply to opt to keep using their old account perhaps, though this might have some ramifications for provisioning.

Implications of LifeLong Learning Support	
Pattern of Use	*Technical*
	Students might need a GUID within MyPortfolio which extends beyond institution boundaries. The NSN is a candidate if MOE business rules permit this.

2.10 Interoperability Specifications and LEAP2A

One of the key concerns expressed by many stakeholders was related to the use of some kind of standard for the XML encoding and transport of portfolio information. The dominant ePortfolio specification (also commonly regarded as the simplest to implement) is LEAP2A. (http://wiki.leapspecs.org/2A/specification). The LEAP2A specification has been implemented in Mahara since version 1.2 for import and export of either a full portfolio or a single view. This is primarily for extracting ePortfolio information for import into another ePortfolio system and vice versa. It is not a mechanism for automated transfer or for more granular information such as a single blog post. Nevertheless the Mahara team have considerable experience with LEAP2A implementation. Many stakeholders expressed the notion that export should use a standard format with LEAP2A as the principal candidate.

3. Scenarios of use

In this section we outline a number of scenarios in which interoperability is required. The technical details in these scenarios have been deliberately left vague or are simplified in order to avoid over-complicating the user story and to resist writing a technical solution into the scenario when alternative approaches are possible to achieve the same ends and the implementation details are yet to be decided.

3.1 Scenario 1: simple scenario outlining export of some structured content from the LMS to the ePortfolio

Student A, in Year 12, is working in her LMS system in which she has just received excellent feedback on her assignment from her art teacher. The assignment consists of a web-page containing a digital image of her latest piece of digital art with her own commentary on the meaning of the piece and the rationale for the approach and techniques she has used. The teacher's feedback is on a separate web-page accessed through a comment link at the bottom of the page. She wants to save it to her ePortfolio in MyPortfolio, so she clicks on the “Save to MyPortfolio” link which is in her sidebar. The LMS packages the required items (the content page text, the photograph and the comment page text) for transport. The LMS contacts the myPortfolio server with a request to post the content to Student A's space. MyPortfolio checks the credentials contained in the request, accepts the content and then sends a response that the content transfer has been successful. MyPortfolio then unpacks the content, works out where to put the resources and then prepares the content for display next time Student A logs in to MyPortfolio herself.

3.2 Scenario 2: a scenario in which a MyPortfolio view is submitted for assessment to the LMS

Student B, in Year 11 has completed a 15 week science research project on water quality in the local waterways. He has gathered and documented all his evidence in his MyPortfolio he has provided a view to stakeholders (including interviewees, experts and his teacher) for comment and feedback. Finally he converts into a view for assessment purposes and wishes to submit it into the teacher's space in the school LMS system. He logs-in to the LMS and sends MyPortfolio a request to export the assessment view to the teacher's space in the LMS. Optionally a mechanism (e.g. xml-signatures) could be employed to ensure that the export is not tampered with after submission.

3.3 Scenario 3: a scenario in which a MyPortfolio view is updated and the teacher and parent are both notified of the change via the LMS and Parent Portal respectively

Student C, a Year 8 at intermediate school has just added a video of her playing her Grade 4 flute piece, flawlessly (she's a flawless flautist) to her ePortfolio and then assigned it to a view for her music teacher and her parents to view. MyPortfolio sends notification messages to both the LMS for the teacher and to the Parent Portal. The notifications simply contain the name of Student C, the name of the view, and a URI to the view itself. Her parents can access the view content after being authenticated in the Parent Portal environment to provide security.

4. Functional Requirements and Architectural Principles

4.1 Functional Requirements

From the concerns and scenarios developed thus far we can derive a set of functional requirements for the interoperability work. These are listed below:

1. Students can export content to MyPortfolio from the LMS.
2. Students can export content to their own space in the LMS from MyPortfolio
3. Students can export a view from MyPortfolio to the teacher's space within the LMS for assessment
4. Exported content should be displayed in the recipient system in a way that remains as faithful as possible to the original content structure The integrity of hyperlinks and original structuring information should be maintained. The receiving system should degrade the content gracefully if it is not possible to display it.[5]
5. Notifications of updates to MyPortfolio are provided to the Teacher in the LMS environment and to the Parents in the Parent Portal
6. Parents can access the content of a MyPortfolio they have been assigned through the Parent Portal
7. Single Sign On to MyPortfolio is provided from the LMS
8. Student accounts in MyPortfolio can be provisioned from the LMS
9. A mechanism is provided for versioning exported content whilst still allowing it to be appended with comments.

4.2 Architectural Principles

We can also draw out a set of general architectural principles from the stakeholder consultations. The first two are really goals which have a bearing on the architecture rather than principles:

1. The architecture should work for all of the three LMS systems (Moodle, Ultranet, and KnowledgeNet) involved in the MLE programme as a reasonable test that it will generalise to a wide range of current and future LMS systems used in New Zealand schools.

2. The architecture should work for a range of school contexts and levels (Primary through to Senior High School).

3. The architecture should use open standards and protocols for the transport layer (e.g. HTTP, XML, LEAP2A)

5 By degrade gracefully, I mean the import process should not fail because some entries cannot be displayed . The system should insert a marker to alert the user that some content didn't export or a similar strategy.

4. The architecture should be consistent with the Common Services Framework (Piers Harding, 2009)

5. Systems should be loosely coupled.

5. LEAP2A should be used as the preferred document encoding for portfolio content transfer

6. The architecture should be a lightweight web services approach. A RESTful approach is recommended over RPC due to its relative simplicity and transparency for service consumers

7. The architecture should use a delegated authentication and authorisation framework for secure transfer of content over HTTP or HTTPS. OAuth is recommended.

8. Attention should be paid to the structure of URIs used for content items such that they are consistent, predictable and meaningful. CURIES should be used.

9. Students should have a unique identifier outside of any particular institution. The NSN is an natural candidate.

5. Architecture Proposal

In this section I propose a high-level model of a web services architecture to support the scenarios and requirements developed so far. The architecture being proposed consists of:

1. A web-services model using a RESTful approach in which MyPortfolio acts as the service provider and the LMS systems act as service clients
2. Content Exchange is handled using the LEAP2A specification
3. Notfications utilise a lightweight ATOM feed
4. The service-based authentication and authorisation model uses OAuth

These aspects of the architecture are discussed in more detail below in order to expose their underlying rationale.

5.1 Web Services Design (REST vs RPC)

The systems under consideration here are web-based systems, which means they communicate over HTTP using the HTTP methods (GET, PUT, POST etc). Resources are encoded using a machine-interpretable representation format(XML, XHTML, MIME, LEAP2A etc) and are located and identified using a URI. These are the basic elements of web-based data exchange.

In any web-based interaction there is a Client or Service Consumer and a Server or Service Provider. In normal use of the web we think of the client as a web-browser being manipulated by a human user to fetch a webpage from a server using the HTTP GET command. The user could also input some information into an HTML form (such as a username and password) and submit it to the server using HTTP POST whereupon the server can do some processing on the information and send a response.

Web services allow both human users and automated processes to do more interesting things than simply fetching static web-pages for display on a client web-browser. In particular they allow a web-based program acting as a client to interact in sophisticated ways with a remote system acting as a service provider. The web service provider exposes an API which allows the client system to request data and perform operations on the remote system.

There are different approaches to designing the style of interactions between service consumer and service provider. The two main alternatives are a REST-type (known as RESTful) approach and an RPC (Remote Procedure Call) approach. In the web-programming world there are adherents to each approach which has been the cause of much heated debate. For the purposes of this paper, the key difference lies in how it affects the design of the APIs for content exchange between MyPortfolio and LMS systems. Using a RESTful approach, the API would consist of URIs to resources (objects described by nouns) on the remote system combined with one of the six operations supported by HTTP (of which GET, PUT, POST and DELETE are the most common). An RPC-style approach involves defining a vocabulary of method calls (such as doSearch, or getList..) for the operations to be performed by the service provider.

In order to understand the differences between these two approaches to web-services design, it helps to look at a concrete example. Amazon's S3 (Simple Storage Service) serves as a particularly useful example. The service essentially provides online data storage on Amazon's servers which you can use as 'cloud' backup for your own servers or as a data host which other people can access. General information about S3 can be accessed here:

http://en.wikipedia.org/wiki/Amazon_S3

The link to the S3 homepage is here:

http://aws.amazon.com/s3/

S3 has two main components: Objects – which are named pieces of data with some associated metadata, and Buckets – which are containers for objects.

A user of the service can remotely get a list of their buckets, create buckets, add objects to buckets and retrieve objects from specified buckets. There are several things that make the approach Amazon have taken in the design of this service interesting from the point of view of this project:

1. It is a web service that provides a similar basic set of operations to those that will characterise content exchange between LMS and ePortfolio systems here
2. It provides A RESTful API to the service.: http://docs.amazonwebservices.com/AmazonS3/2006-03-01/
3. It also provides an RPC style API wrapped in SOAP http://docs.amazonwebservices.com/AmazonS3/2006-03-01/

This makes it very easy to understand and compare the two API styles

4. S3 defines an authentication model which relies on public and private key encryption. Where the amazon service uses its own custom framework for this, it is proposed here that the OAuth framework would perform a similar function.

Whilst the choice of approach (REST vs RPC) may ultimately impact little or not at all on the end user, it is an important decision for the developers of both the APIs and the client code in participating systems and is worthy of early discussion. This (over-simplified) introduction is to provide a context for the discussion which doesn't preclude non-technical stakeholders.

5.2 Which system is the Service Provider and which system is the Client?

Since we have identified the need for bi-directional communication between LMS and MyPortfolio systems for content exchange, there are two possibilities for the overall architecture.

In the first scheme APIs could be implemented on both sides and the system initiating the export request could be regarded as the client (making a request to the remote system to PUT some information in the user's space there and sending the users authentication details with the request). The remote system (depending on where the request originates) is the service provider, exposing an API which is the target of the request.

In the second scheme an API is only implemented in MyPortfolio, i.e. it always acts as the service provider in exchanges with LMS systems with which it is associated. The service requests always originate from the LMS (using either a PUT action to export content or a GET action to retrieve some content) .

The two schemes are represented diagrammatically in Appendix 2.

It is quite obvious for a host of reasons that the second approach is preferable:

1. API development and the associated processing code only needs to happen in MyPortfolio. LMS developers can confine their efforts to writing client-side code which implements the service calls to the API published by MyPortfolio and processing any content received.

2. Since there are many LMS systems and only one MyPortfolio service this scheme instantiates a many-to-one model instead of a much more complex and inefficient many-to-many model that would result from scheme 1.
3. This same general model (scheme 2) also provides an effective approach for provision of notifications about updated content to service consumers (LMS systems / Parent Portals). MyPortfolio The service consumer would be required
4. Using this model and an appropriate authentication mechanism using public and private key encryption, it is possible that teachers could be provided with access to pull certain content from a students portfolio where it has been identified as exportable and the student has provided the teacher with authorisation and a signed URL to access it. The teacher then has a snapshot of the portfolio work for purposes such as: moderation, assessment and monitoring of progress.

5.3 Mahara, Moodle and MNet

http://docs.moodle.org/en/Development:Mahara_Portfolio_Plugin

http://docs.moodle.org/en/Development:Portfolio_API

http://docs.moodle.org/en/Development:MNET_API

The Portfolio API developed in Moodle 2 to allow export of content to Mahara uses MNet (a set of XML-RPC libraries originally designed to allow interoperability between Moodle instances) to export content to Mahara. MNet includes functions for establishing connections between servers, handling authentication and authorisation – and through the Portfolio API, exchanging content. This begs the question, why not use MNet and the Portfolio API as an architecture to support the more general LMS-MyPortfolio work being proposed here?

Whilst undoubtedly the MNet model provides a very useful technical reference for this development and deserves detailed consideration, there are a number of issues with Met:

1. Some of the LMS systems (e.g. Ultranet and KnowledgeNet are proprietary products, under the licensing scheme for MNet as part of an open-source project it cannot be used for proprietary products.

2. At present (to my knowledge) LEAP2A export is not implemented in this scheme though there are plans to extend it to accommodate LEAP2A in the future. The current model exports content as single files (a blog post), thus structuring information around complex content is lost.

3. MNet was developed to support a different scenario than the one being proposed here, thus it may prove to be more complex and error-prone to attempt to adapt it (even if the licensing scheme permitted) than develop a clean API from scratch.

5.4 Content Exchange using LEAP2A

What is LEAP2A

http://wiki.leapspecs.org/2A/specification

http://datatracker.ietf.org/wg/atompub/charter/

LEAP2A is an XML specification for the transfer of ePortfolio Information between systems. It is built on the Atom Publishing Protocol (APP). Atom is a feed format (similar to RSS) for representing published web resources such as Blogs, Journal Articles etc. It provides many of

the basic structural elements for representing ePortfolio information. LEAP2A extends and specialises Atom by providing additional elements for specific ePortfolio related information and interpreting certain Atom elements in a specific way.

The Structure of LEAP2A

A full description of LEAP2A elements is available in the specification. Here it is worth noting a few basic points.

The top level element in a LEAP2A document is the atom:feed element. Each feed element has to have as mandatory elements:

an id – a URI which uniquely identifies the resource

a Title - anything

an updated element – a timestamp of export

If the author element is also included at feed level then it propagates to each entry.

Individual portfolio items are represented as atom:entry elements. Each entry has its own unique identifier (URI). Since Atom is basically designed to handle blog posts, it doesn't distinguish types of entry elements, so LEAP2A extends Atom by providing a set of portfolio entry types with attributes specific to that entry type.

Entries define relationships with other pieces of content (themselves entries) in order to represent structured content. LEAP2A extends Atom by defining a set of relationship types.

Where a portfolio item is associated with a file (e.g a JPEG image, or a PDF file) that cannot be represented as an entry, the enclosure relationship type is used to link the entry to the file itself.

Once the XML document has been constructed, it is saved as: leap2a.xml and is packaged in a zip file with any files that have been referenced within the document. The zip file is then ready for export.

Why use LEAP2A

- It provides a standard way of exchanging ePortfolio content as an XML document
- It is the best-supported and simplest to implement of the portfolio specifications available
- It is capable of representing structured content and handling external files
- It represents person and organisation information as well as specific portfolio content related information
- It is based on a robust format (ATOM)
- It is open and extensible

5.5 Notifications using an ATOM feed

For the Notifications mechanism, I have proposed using an ATOM feed. This is a natural and obvious suggestion since this kind of automated notification requirement is exactly what web feed formats were designed to handle. Whilst RSS 2.0 could be used instead of ATOM 1.0 for this purpose, there are several distinct advantages to using ATOM. The main properties of the two specifications are discussed in the following web articles:

http://www.intertwingly.net/wiki/pie/Rss20AndAtom10Compared

http://www.problogger.net/archives/2006/03/30/rss-vs-atom-whats-the-big-deal/

We have already described the basic structure of an ATOM feed in the context of the LEAP2A specification. As has already been noted, it is desirable to keep the payload of the notification as light weight as possible for reasons of minimising the security risk and load on the service provider's server. Thus the ATOm fields the notifications might employ would be just the mandatory fields:

- ID – The URI of the updated view in MyPortfolio
- Title – The name of the view
- UPDATED – A timestamp of the last update to the view

In addition the name of the student would be very useful for teachers:

- AUTHOR – The name of the student

5.6 OAuth as an Authentication and Authorisation Framework

http://oauth.net/

The web-services model being proposed here requires server-to-server transmission of the students' protected, private that is stored either in MyPortfolio or in the LMS. Normally if the student wished to access her data stored in one of these systems, then she would have to log-in to the system using her credentials (username and password) in order to gain access.

In the interoperability scenario being proposed here the student is making the request to her MyPortfolio account via another website (the LMS). The LMS needs some credentials in order to access the data in her account, but simply allowing it to pass her username and password as part of the payload of the HTTP Request would be an unacceptable security risk.

OAuth provides a system of token exchange (similar to public and private key encryption) which can be used over HTTP or HTTPS to give much higher security in this kind of server-to-server interaction using a web-services model.

6. Recommendations for Further Work

This paper proposes a high-level model for a technical architecture to support interoperability between LMS systems commonly used in New Zealand schools and the Mahara ePortfolio platform being offered as a nationally hosted service to New Zealand schools under the name of MyPortfolio http://myportfolio.school.nz/ . The model is designed to support a variety of stakeholder concerns collected from the sector during the course of this work. So what are the next steps from here? In this final section of the report I briefly outline the work that will be required.

6.1 Wider community exposure to the report for feedback and participative input

This report represents a start in the process of stakeholder consultation. To facilitate further collection of stakeholder viewpoints, a Wiki-Educator project has been created. To provide feedback on this report and have your input into the process, please visit the Wikieducator pages: http://wikieducator.org/LMS-MyPortfolio_Interoperability_Project

6.2 API development in MyPortfolio

As indicated in this report, the responsibility for development of the web-service API to support the interoperability described here lies with Mahara developers. Mahara is an open source project, thus it is essential to liaise closely with the Mahara development community to bring them on board with the project and mostly importantly to guide the development of the web-services

6.3 Development of service clients within LMS systems

The service clients will be developed by those responsible for the LMS systems. Close communication is going to be required, between the LMS developers and the Mahara development team working on this project. Appropriate project management structures will need to be put into place and funding made available to support the work.

6.4 MyPortfolio governance, support and development of training materials

This project is built on an assumption that MyPortfolio will continue to be provided to New Zealand schools as a national hosted service. Mechanisms need to be put into place to ensure ongoing funding for the governance of MyPortfolio, the provision of support as required and the development of training materials and guidance for schools.

7. Supporting Material

7.1 Appendix 1: Acknowledgements

I would like to thank all those who have taken the time to provide input into the project and have given invaluable feedback on initial drafts of this report. Those people include (but are not limited to) the following

1. Mark Osborne, Albany Senior High School
2. Sarah Martin and Chris from Stonefield Primary School
3. Gerard McManus, Papatoetoe
4. Hugh Nettar and colleagues at Ultranet
5. Dave Braddock of Dataview
6. Kristina Hoeppner, Mahara / Catalyst IT
7. Ian Munro and Paul Seiler, Ministry of Education
8. Yuliya Bozhko, Massey University
9. Carolyn Stuart, Tawa intermediate School.
10. Piers Harding, Catalyst IT

7.2 Appendix 2

Diagrams depicting possible architecture schemes for LMS-MyPortfolio content exchange.

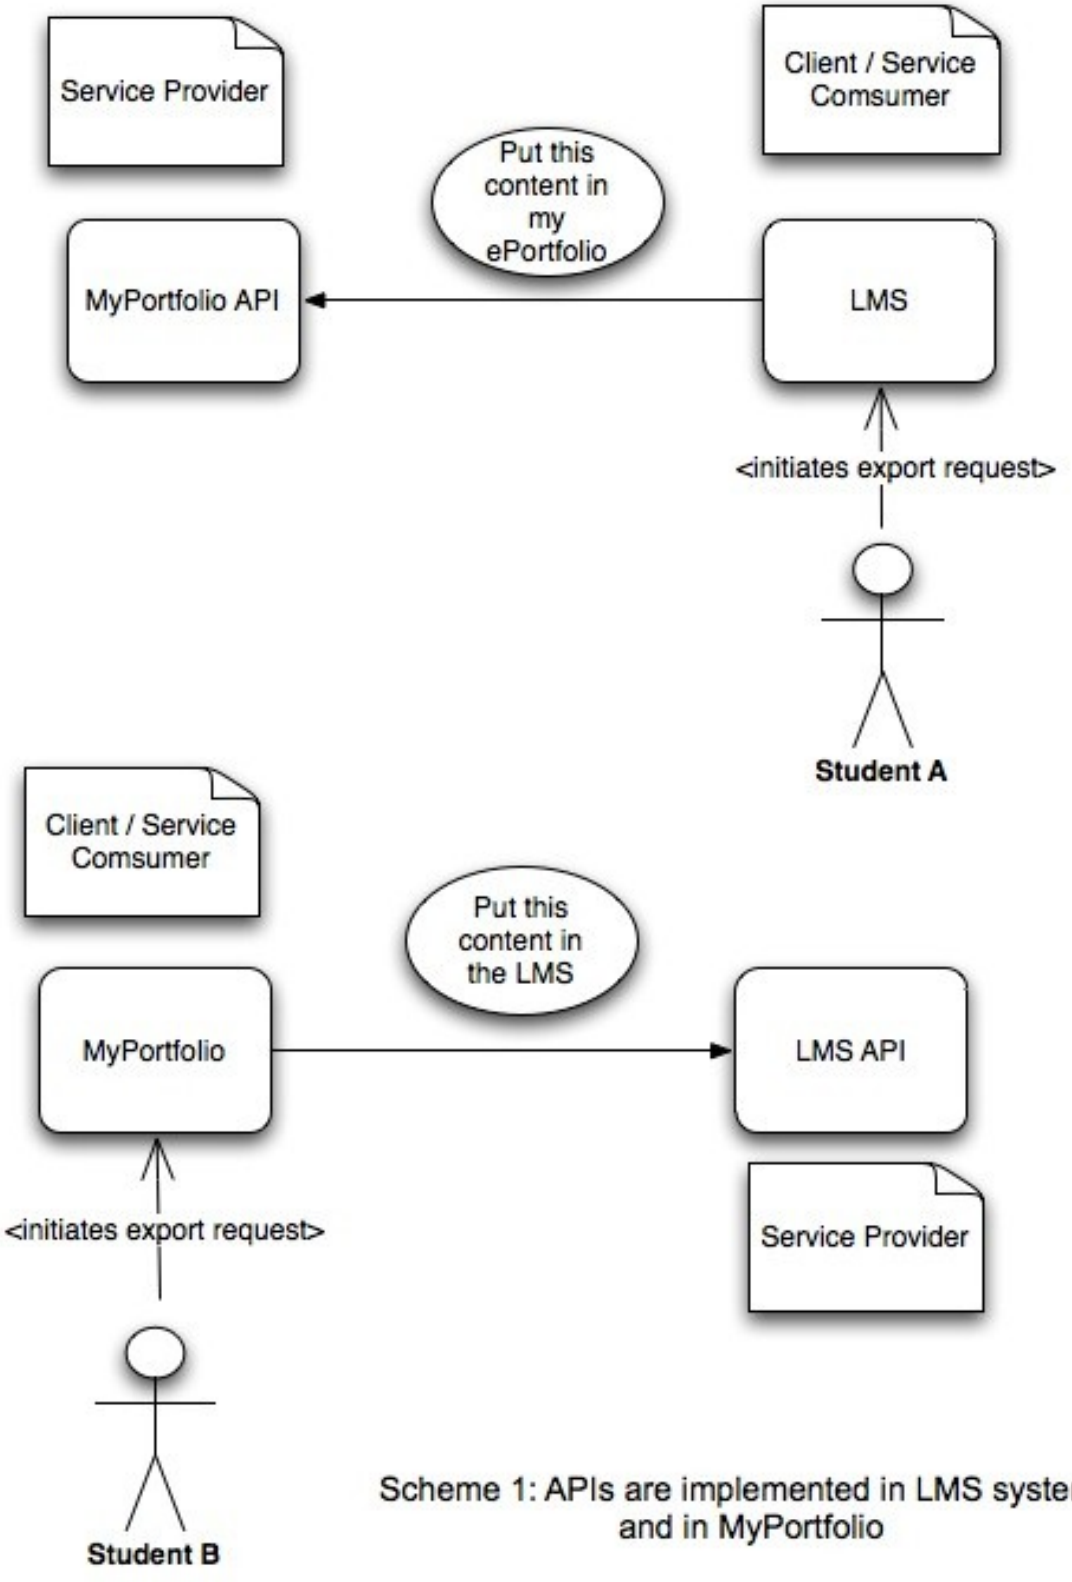

Scheme 1: APIs are implemented in LMS systems and in MyPortfolio

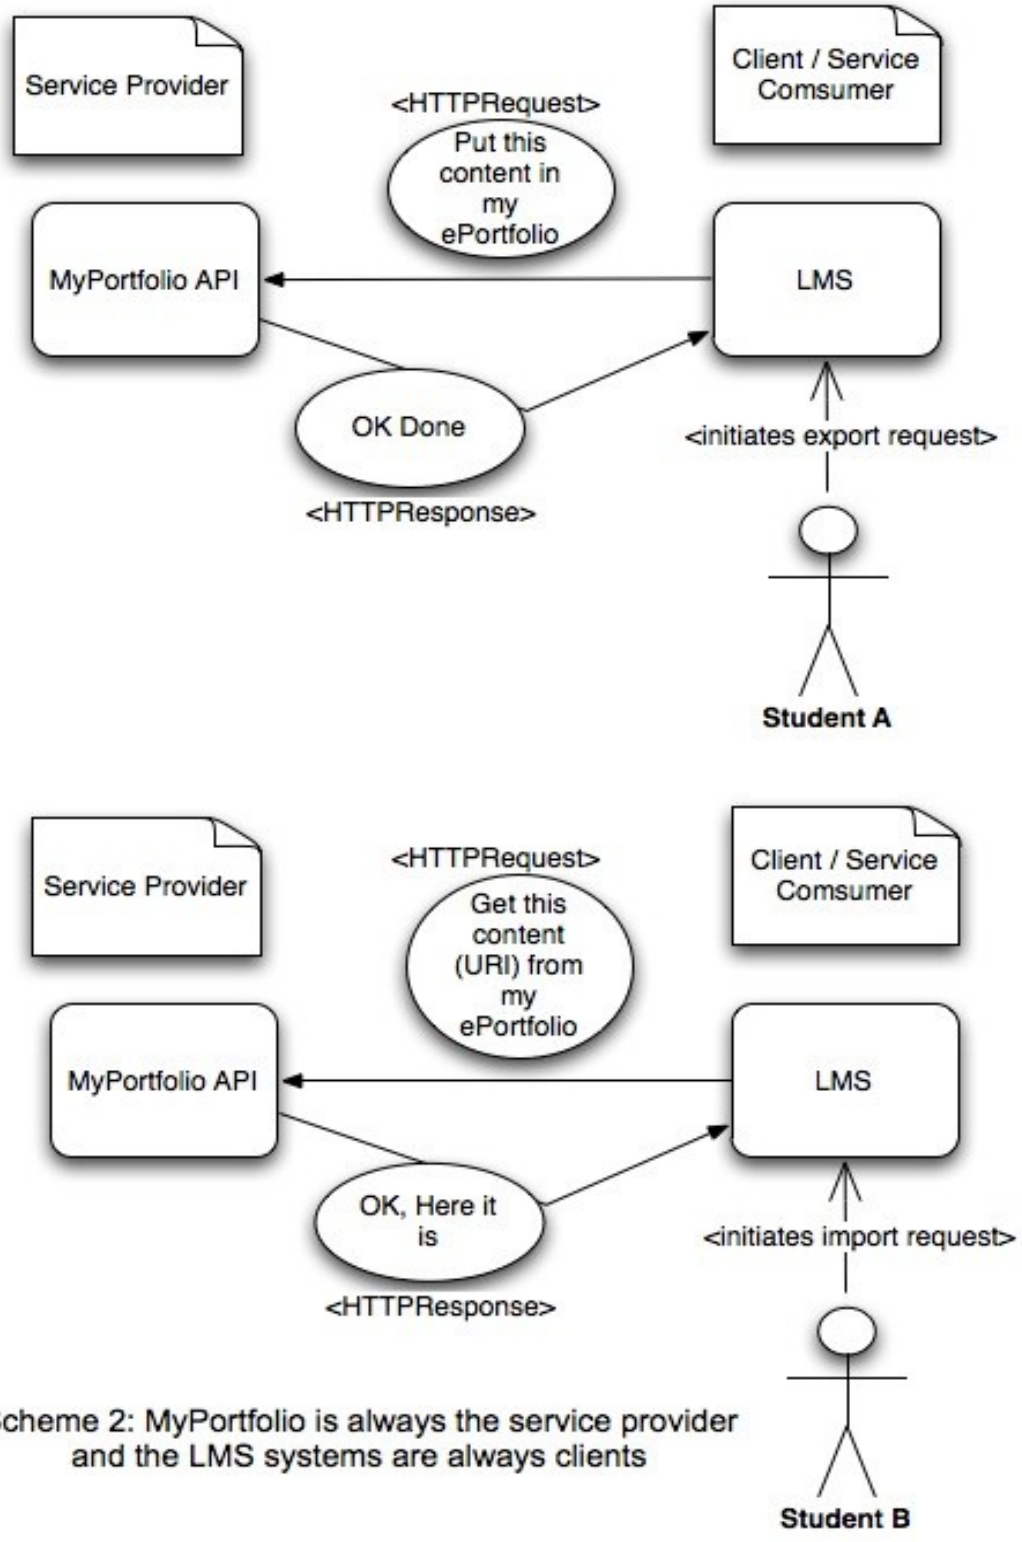

Scheme 2: MyPortfolio is always the service provider and the LMS systems are always clients