




Talis White Paper:
Project Silkworm.



In the age of
Google,
Amazon
and **MSN,**
why is content
in the library
domain still so
difficult and
expensive to
discover,
access and
share?



	Overview	3
1:	The experience of the Google generation	4
1.1	What do users of Internet services experience?	4
1.2	What do users of library systems experience?	5
1.3	A comparison: Amazon.com versus the OPAC	5
1.3.1	The Amazon website experience	5
1.3.2	The Library OPAC experience	6
1.4	The need for Collaboration	7
1.4.1	A radical change of perspective	7
2:	The limitations of technology today	7
2.1	The Web 1.0	7
2.2	The Hidden Web	8
2.3	Islands of content	8
2.4	Prohibitive costs	8
3:	The possibilities for technology tomorrow	9
3.1	The Web 2.0	9
3.2	A fundamental change in the technology stack	9
3.3	Amazon web services	10
3.4	Web 2.0: Hype or reality?	10
4:	The impact of Web 2.0 on the library & information sector	10
4.1	The possibilities Web 2.0 offers the library and information domain	10
4.2	A shared platform	10
4.3	The need for collaboration between competitors	10
4.4	The need for an architecture of participation	10
5:	Making it happen: Project Silkworm	11
5.1	Introducing Project Silkworm	11
5.2	The values of Project Silkworm	11
5.3	What is Silkworm?	12
6:	The Silkworm platform	12
6.1	Silkworm is a web 2.0 platform	12
6.2	Silkworm's Platform Architecture: Nodes	12
6.2.1	Nodes, Networks and Peer to Peer	12
6.2.2	Network effect or Metcalf's law	13
6.2.3	What are the nodes in the Silkworm network?	13
6.2.3.1	Content Nodes	13
6.2.3.2	Application Nodes	13
6.2.3.3	Aggregation Nodes	13
6.3	Silkworm Platform Architecture: Services	13
6.3.1	Directory services	14
6.3.2	Access services	14
6.3.3	Aggregation and sharing services – BigFoot	14
6.3.4	Grouping services	15
6.3.5	Notification services	15
6.3.6	Transformation services	15
6.4	Silkworm's Platform Architecture: Participation	15
6.4.1	he need for openness and an architecture of participation	15
6.4.1.1	How to share content: Creative Commons	15
6.4.1.2	How to share code: Open Source	16
6.4.1.3	How to share contracts: Open Standards	16
6.4.2	Bigfoot aggregation services	16
6.4.2.1	Removing the walls between application instances	16
6.4.2.2	Removing the walls between different applications	16
6.4.2.3	Bigfoot – Participation without walls	16
	Conclusion	17
	About Talis	17
	About the Author	17

Overview

Currently, users of Internet services such as Google, Amazon, eBay and Yahoo are enjoying fantastic, participative experiences.

Unfortunately, the same can't be said about users of traditional library and information systems. Here, the problem of silos of content held in different formats and hidden to a wider audience is preventing vendors from easily and cost-effectively building systems that can provide users with expansive and rich experiences.

This white paper discusses these issues and introduces Project Silkworm: Talis' response to the challenges facing library and information systems vendors. This research project has investigated ways to make the experience of library users akin to those enjoyed by other Internet services.

As this paper will explain, Project Silkworm is based on the concept that library vendors must now collaborate in order to begin to deliver better services. This focus on participation (of both vendors and users) permeates the whole project and is captured in four key values:

- Sharing and community over duplication and isolation
- Reuse over reinvention
- Openness and interoperability over exclusivity
- Experimentation over certainty

The project centres on the creation of Silkworm, a Web 2.0 service-orientated network platform for discovering, sharing and consuming content that is currently hidden from the web.

The paper will discuss technical concepts but it is not just a technical paper. Project Silkworm proposes a fundamental shift in the way the library and information domain operates and, consequently, this paper covers business, economics and technology. A clear view of the bigger picture is vital, Talis believes, to understand the forces that will drive this change.

This paper may appear radical at first. However, reflection on the speed and depth of change that technology is having on modern life will reveal a simple truth: the pace of change today is staggering. It will soon be very difficult to describe to children how we used to find information and communicate before mobile phones and the Internet. Yet a decade ago the Internet and mobile phones weren't part of our everyday lives. This revolution has happened in a few short years.

Consequently, the Silkworm platform is already under construction and will deliver value immediately through the close collaboration of the Silkworm development team with various partners.

For more information and ways of getting involved visit <http://silkworm.talis.com>

1: The experience of the Google generation

Today, when it comes to finding content and information, people are enjoying a fantastic experience. With organisations such as Google, Amazon, eBay or Yahoo leading the way, users of these sites have come to expect that they will be able to find resources quickly and with minimum effort.

1.1 What do users of Internet services experience?

Users of the Internet services offered by Google, Amazon, eBay and Yahoo have come to expect an experience that is:

- All Encompassing
- Without Walls
- Immediate
- Rich
- Participative
- Personalised
- Fulfilling

All Encompassing

With Google, as far as the user can tell, you are searching all web sites; Amazon has all in-print books; and eBay provides access to a multitude of auctions. Quite simply, there is no need to go elsewhere.

Without Walls

There are no frustrating barriers as you progress through these sites either. Using rich navigation, you are able to move around the information space from one item to another easily. You don't keep bumping into walls and having to change system, search and navigation. For example, before Amazon existed, the purchase of new and second-hand items usually took place in separate shopping environments. Now everything is in one place – including the ability to purchase the books.

Immediate

For the most part, users of these sites are satisfied quickly and easily. Anything that the user wants is rarely found more than one or two clicks away at most. The systems are fast and accurate. Plus, there are often multiple ways to help the user locate whatever it is they need as quickly as possible.

Rich

No effort is spared to bring all relevant information and content together to help users achieve their goals. Reviews, book jackets, author interviews, plus serendipity mechanisms like Amazon recommendations, make the experience anything but dull.

Participative

If desired, the user can get involved with these sites through user reviews, supplier ratings, discussion groups or reading lists. eBay, for example, is built almost entirely on participation. Users build the auctions, bid in the auctions and leave feedback afterwards. All this work then feeds back into the rich experience for others and increases the value for all. The result is a strong network effect.

Personalised

These sites often use specific and intelligent personalisation. For example, Amazon suggests books to users based on what they have already purchased.

Fulfilling

The systems used by these companies ensure that the user's request can be fulfilled as quickly as possible. Amazon shipment, Google linking, and eBay payment mechanisms are all focussed on supplying satisfaction every time. This is a reflection on the strong gravitational pull that users now exert on services – expecting goods to come to them via home delivery service options.

1.2 What do users of library systems experience?

Unfortunately, however, users simply don't enjoy the same fantastic experience with library and information systems.

Consider, for example, what happens when people want to find a book to borrow. If library users know the exact book they want to read then they may perform a search on the library OPAC. If they are lucky, the library will stock the book and they will be able to visit the library, borrow the book or reserve it if it's on loan. With some libraries, they may even be able to reserve the book online. In all cases, though, in order to borrow the book, the user must go to the library.

If the book is not in stock, however, the search on the OPAC will bring up no hits. The user will not even be able to tell whether that's because the library doesn't stock it or because the book doesn't exist.

The system is even more unhelpful if users don't know what they are looking for. Library OPACs provide only the barest of details and offer no help if they are searching for fiction or have to decide between several books.

These are just some examples of how the traditional library and information systems fail to deliver a fantastic and engaging experience. There are many more.

1.3 A comparison: Amazon.com versus the OPAC

The extent of the gap between the experience of the library service and that of services on the Internet becomes clear through a comparison between using Amazon.com and using a typical library OPAC.

1.3.1 The Amazon website experience

Users of Amazon enjoy an engaging experience:

All Encompassing

Amazon provides a one-stop-shop for print material and second hand books. It's rare that you can't find what you want.

Without Walls

There are many links to all sorts of related content on the Amazon website. You have access to author links, subject links, related books, reviews and recommendations – without you needing to visit anywhere else.

Immediate

Everything that you need is only one click away including your own participation.

Rich

The information available to you is comprehensive. You can see Bibliographic and holdings details, images, book reviews, editorial reviews, search inside the book, wish lists, and useful data such as 'people who bought this book also bought these items' (see screen shot).

Participation

You are encouraged to get involved at all stages through invitations to write reviews and guides, requests for feedback and ratings, or the ability to be added to a list.

Personalised

You can easily see what you have looked at so far today, your wish list and your suggested reading (not shown in the screen grab).

Fulfilling

Amazon offers multiple methods of fulfilment with exact information on availability and dates. Often, users can expect next day delivery.

1.3.2 The Library OPAC experience

Users of the Library OPAC, on the other hand, have a fairly limited experience.

All Encompassing

Users can see only what exists in this library. If they want to find an item that isn't stocked they must search elsewhere. There's no way of knowing whether an item is missing because the library doesn't stock it or because it simply doesn't exist.

Without Walls

The author and subject information is only helpful if the library has the related works in stock. But, if users want to find other related information for this work they must go to another system and search.



A screen shot of the book details screen on Amazon

Immediate

Unlike the Amazon example, the library OPAC doesn't provide options where you can obtain the book immediately.

Rich

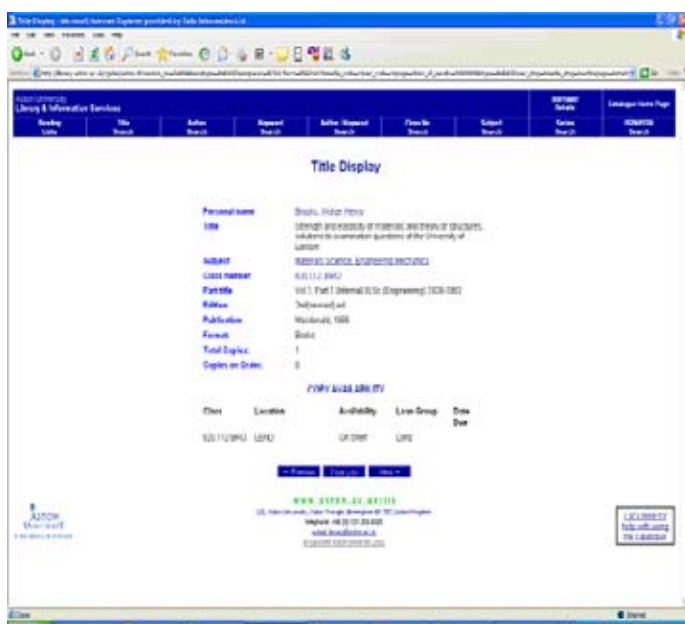
The library OPAC only has bibliographic details and holding information.

Participation

There is literally no way to participate in the Library OPAC system.

Personalised

The standard page is completely standard with no options for providing personalised service.



A view of a typical Library OPAC

Fulfilling

The library OPAC has some information that helps users to know whether or not it's worth visiting the library. However, this is the only method fulfilment provided: there are no options for home delivery etc.

1.4 The need for Collaboration

1.4.1 A radical change of perspective

Nearly every person in the UK has had some experience of Google, Amazon, eBay or Yahoo, and many of them are now beginning to demand that these experiences are replicated in other areas. This is a benchmark that's only going to get higher as these companies expand their reach into more aspects of people's lives.

So how can library and information systems begin to deliver on these expectations?

Talis has spent a considerable amount of time investigating ways that libraries could improve the experiences they offer users. To do, so successfully, we believe, requires a radical change of perspective.

- The library is not an end in itself

The users' world is much larger than the library world. The library is but one means of fulfilling their needs. If the traditional information services don't match up to the experience that the user wants, then users will and are going elsewhere.

- The technical and organisational challenges for institutions and vendors will need to be solved in a collaborative way.

One of the main challenges is removing walls or silos and this requires collaboration especially between competitors. Institutions and vendors will all benefit from sharing in some areas and competing in others. Standards, open source software and creative commons give us unprecedented ways to collaborate and yet still compete.

- A shared platform is required to lower costs, break down walls and provide an all encompassing experience for users.

The market is structured in a high cost way with many vertical vendors and little horizontal specialisation. Any mature market needs horizontal structure to lower the total costs for all. There are technology pieces missing but it is too expensive for each player to provide an all encompassing experience, the market should not have to bear the weight of each vendor duplicating this.

- There is massive value to be unlocked by harnessing the efforts of the users across the industry with a shared architecture of participation.

2: The limitations of technology today

"Why does content in the library domain remain so difficult and expensive to discover, access and share?"

Why can't we Google the library domain in the same way we can Google the web? There are reasons why this is currently difficult.

2.1 The Web 1.0

The World Wide Web grew up on a foundation of widely shared standards (mainly HTTP, HTML) and was built upon a rich interlinking mechanism (the url). So the web is standards based and densely interconnected. The pages were designed for humans to navigate and so the content was largely unstructured.

However, it quickly became apparent that search engines could discover nearly every page on the Internet by following the links from page to page. Search engines therefore built huge free text indexes so that users could search and find what they needed.

Unfortunately, though, there was little standardised metadata to assign any meaning to content. That meant that, as the amount of content grew, finding what you wanted became harder.

Google and other search engines realised that by analysing the link structures themselves some weighting of value could be added to return more accurate results. While this is undoubtedly an improvement, Web 1.0 remains human readable with machines unable to understand the data in most web pages.

2.2 The Hidden Web

The vast majority of the world's information is on the Internet (although not always the web). Often, it is held in databases that may or may not have a web-based interface but, undoubtedly, the contents of the databases are not in HTML and are not discoverable by following URL links.

That means that Google and other search engines cannot find this content. This raft of information that is available but cannot be seen or searched easily has become known as 'the hidden web'.

This content has strong metadata and structure and can be understood by machines but is not standards based and not discoverable.

2.3 Islands of content

Library catalogues, holdings information, archives, institutional repositories and specialist databases used by researchers are generally examples of hidden web content.

This is because most traditional library systems grew up before the Internet was generally available and certainly before the World Wide Web existed. The industry took a pragmatic view of the Internet and used it to connect applications together and to provide user interfaces via the web. However, the Internet exists on the outside of the applications. Each system is therefore an island of content.

2.4 Prohibitive costs

Sharing and searching across these islands of content has proved to be a tough problem. Not least because sharing content is potentially very costly.

Talis, for example, created a shared cataloguing application called Talis Base so that the community of Talis libraries could share the effort of cataloguing. UnityWeb was also created to share the holdings information between libraries to overcome the problem of finding a copy of a book for Inter Library Loans (ILL).

However, these systems for sharing content were expensive to build. Partly, this was due to the fact that the systems had to be built from scratch each time because there was no platform to build on and, partly, this was because the systems were proprietary and so could not be easily shared.

3: The possibilities for technology tomorrow

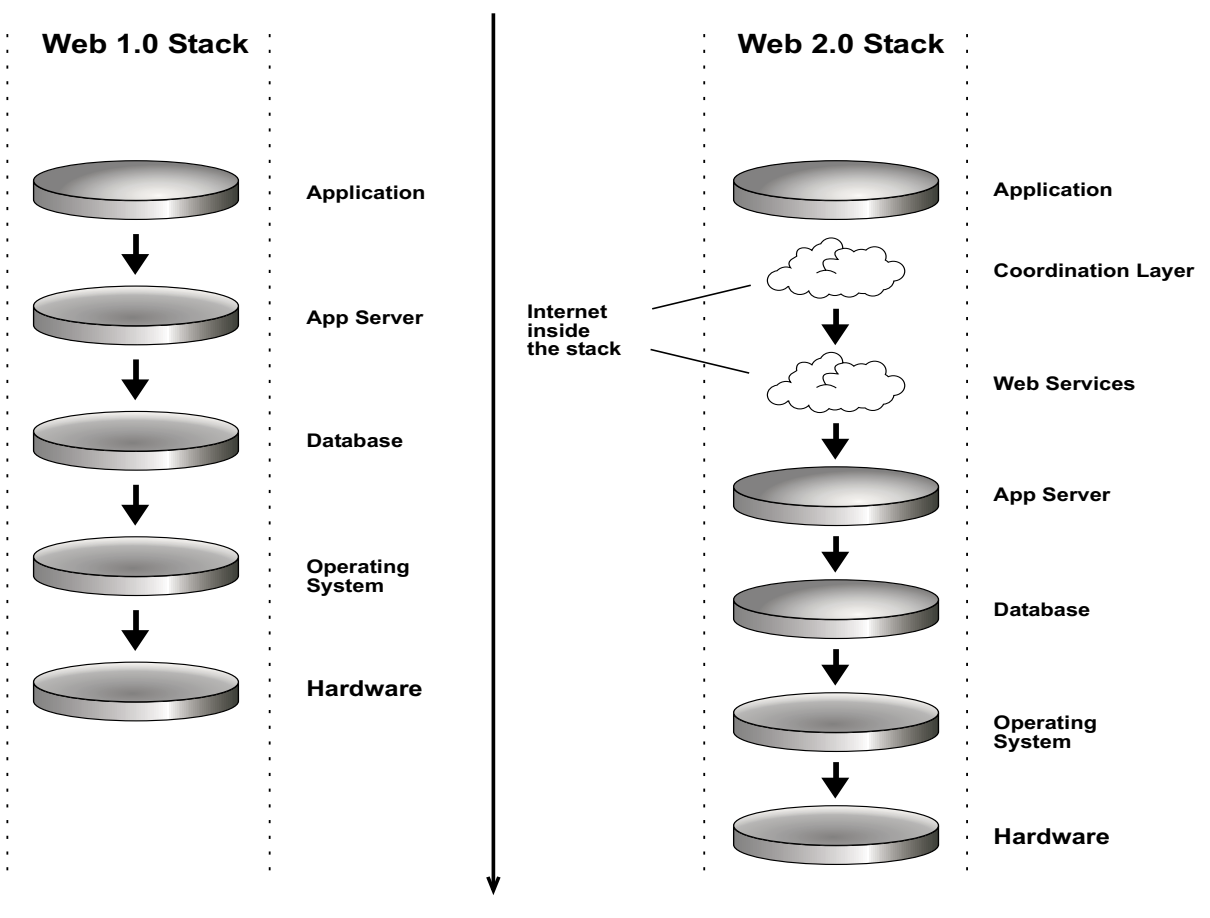
3.1 The Web 2.0

Web 2.0 is the evolution of the Web 1.0 into a network that supports machine as well as human interfaces. Web 2.0 is about machine interfaces being available so that value, be that content or transactions, can be reused and built upon instead of being locked up in silos. In the vision of Web 2.0, applications are assembled from web services that are running somewhere on the Internet.

3.2 A fundamental change in the technology stack

Web 2.0 is a fundamental change in the technology stack that we have been living with for the last 15-20 years. This change is occurring due to universal connectivity and standardisation combined with the massive increase performance/price.

In the early days of the Internet, it was simply bolted onto our existing technology stack. The Internet connected applications together. From the application point of view, the Internet was on the outside.



So more and more, the value of the application is in the assembly of web services. Moving the value of collaborating into the sharable standards based web services is the perfect way for vendors to build applications yet share and enrich a common value, breaking through the silos.

3.3 Amazon web services

There is, perhaps, nowhere where this is more clearly seen than with Amazon web services. Amazon recognised that the content powering the Amazon website was locked into the silo of a human readable HTML web site. The company understood that if the content was available via machine interfaces it could be used by others to create more applications built on Amazon. By enabling this collaboration by others, the reach and value of Amazon would be increased.

3.4 Web 2.0: Hype or reality?

Some people might argue that Web 2.0 is already becoming over-hyped. And, to a degree, that's true. It's a natural part of technology adoption that the hype always exceeds the reality in the early stages (this is known as the hype cycle). However, as the technology matures this ceases to be the case.

For example, as the Internet began to grow many said its importance was becoming over-exaggerated. It's clear today, though, that the reality is in fact exceeding most of the hype. Talis believes that Web 2.0 will fulfil its initial promise. However, such a major change in the technology stack is not going to happen overnight.

4: The impact of Web 2.0 on the library & information sector

4.1 The possibilities Web 2.0 offers the library and information domain

So what does all this mean?

The impact of Web 2.0 on libraries promises to be profound. The advances in technology mean that, at last, libraries have a realistic way of sharing content and providing users with excellent experiences.

4.2 A shared platform

Using the technology of Web 2.0 and service-orientated architectures, software vendors will be able to remove the walls or silos of content and build applications based on a shared infrastructure. Here, the value will be generated by the applications built on this platform and the users that engage with it.

Currently, the library market is structured in a high cost way with many vertical vendors and little horizontal specialisation. It will be too expensive for each player to provide an all-encompassing experience and the market should not have to bear the weight of each vendor duplicating this. Vendors will therefore need to work together. Indeed, it's the sign of any mature market that a horizontal structure is required to lower the total costs for all.

4.3 The need for collaboration between competitors

Perhaps the greatest upheaval for libraries will be the fact that removing walls or silos will require collaboration between organisations that are currently competing. It doesn't mean an end to competition, but institutions and vendors will benefit from sharing in some areas and competing in others.

It's a completely new way of working. However, Talis believes that the industry has the maturity to collaborate and share what is in common and compete where we add unique value.

4.4 The need for an architecture of participation

There is massive value to be unlocked by harnessing the efforts of the users across the industry with a shared architecture of participation.

So what is this?

An application may allow users to get involved and contribute content to the system, if they wish to, in a way that builds value for all other users of that application. Such an application could be said to have an architecture of participation: its construction supports a participation in and enrichment of the user experience by the users themselves. Amazon and eBay are classic examples of this.

Clearly, the value of this user participation increases as numbers of users grow: the more users, the richer the experience for all.

If, as an industry, the library domain can engage its users, then it will be possible to harness the experience and efforts of millions to add value for all. One very simple example is that, if a user wants to write a review whilst using an OPAC from Vendor A, then that review should be available in the OPAC of Vendor B and vice versa.

5: Making it happen: Project Silkworm

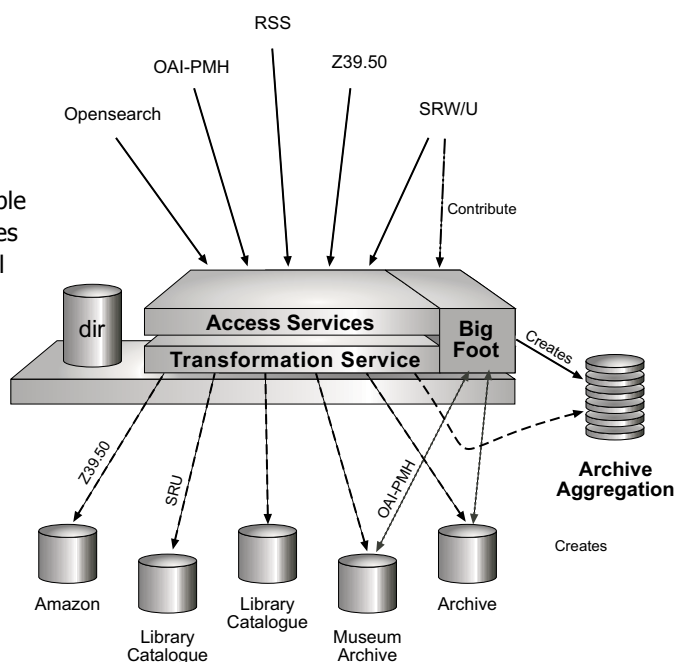
5.1 Introducing Project Silkworm

Project Silkworm is a Talis research project that addresses some of the questions posed in this paper.

The Talis Research Group has spent considerable time investigating possibilities for helping libraries to make content more visible and accessible to all and to give users the same quality experience they currently enjoy elsewhere.

Under the banner of Project Silkworm, Talis has begun to create a Web 2.0 service-orientated network platform designed to discover, share and consume content that is currently hidden from the web.

The platform can be considered as a coordination and access layer sitting above the hidden web that allows hidden web content to participate in the Web 2.0 paradigm and new applications to be built that have powerful architectures of participation to enhance the experience of users.



5.2 The values of Project Silkworm

The key aim of Project Silkworm is to remove the technical and cost barriers to building applications that share and use content. Talis hopes to engage the industry in a shared effort to provide users with the all-encompassing experiences they are demanding.

Therefore, Project Silkworm is founded on the following core values:

- Sharing and community over duplication and isolation
- Reuse over reinvention
- Openness and interoperability over exclusivity
- Experimentation over certainty

These values are very important and set the tone and direction for the whole project. The values of sharing, community, openness and interoperability are not optional; rather they are essential if libraries are to succeed in delivering on the promise of a shared platform. That said, there is much work still to be done. Through experimentation, Talis hopes to support innovation, agility and the chance to try multiple options.

5.3 What is Silkworm?

Silkworm is a service-orientated web 2.0 platform for making hidden web content discoverable, accessible and consumable.

This shared infrastructure provides a platform that application vendors (including Talis) can use to build applications that are not limited by the current silos and benefit from the massive network effect of collaboration. With Silkworm, value is generated by the applications built on it and the users that engage with it.

By definition the vast majority of applications that can be built on Silkworm are not Talis applications. The full richness of Silkworm is available to the whole development community including other Library vendors.

Therefore, Talis believes that Silkworm is one way that the industry will be able to lower the technical and economic cost of building the experiences that our users now expect.

6: The Silkworm platform

6.1 Silkworm is a web 2.0 platform

The Silkworm platform is a service-orientated architecture for making the hidden web discoverable, accessible and consumable. It provides a layer that exposes the hidden web as a web 2.0 platform on which applications can be built.

Like other web 2.0 platforms such as Amazon web services, you don't need to download and install anything. The platform runs in the Internet cloud and the services are available to use from any application or website that has external Internet access.

6.2 Silkworm's Platform Architecture: Nodes

Silkworm is a network platform built on the principle of service orientation and, as such, can be considered as a platform for coordinating and orchestrating the space between services. This is in contrast to the traditional operating system platforms that manage the space between a service interface and the machine hardware on which it runs.

6.2.1 Nodes, Networks and Peer to Peer

As a network platform, Silkworm is quite different from an operating system platform or a database platform. Unlike the others, a network platform is concerned with the emergent qualities of harnessing multiple instances of an application or service. These instances are nodes in a network. It is therefore necessary to combine network theory with software engineering to successfully build network platforms. There is no single location you can point to and call it the network platform. Just as you cannot point to anything and say: that is the Internet.

6.2.2 Network effect or Metcalf's law

One of the most important aspects of networks is described by Metcalf's law. This states: "the value of the network increases as the square of the number of nodes".

A simple example of this is the telephone network. Consider a single handset as a node. If you are the only person in the world with a phone, its value is only as much as you can do with that node on its own. If two people have phones, however, it becomes a little bit more useful, even though the phone is no more useful by itself than before.

As more people have phones the value of owning a phone increases very quickly. Because of the squaring law, the value of the network soon outweighs the value of the node. There are a large number of subtleties that affect the exact scaling law but the guiding principle that the value of the network vastly exceeds the sum of the parts is clear.

6.2.3 What are the nodes in the Silkworm network?

There are three main node types, content nodes, application nodes and aggregation nodes.

6.2.3.1 Content Nodes

This includes any collection of metadata or content that is online and can be searched or harvested via standard protocols such as Z39.50, SRW/U, OAI-PMH etc. This could be archive repositories, library catalogues, specialist databases and so on. The library and information domain is replete with content nodes that can be registered with the Silkworm network.

6.2.3.2 Application Nodes

Application nodes are created by any application that wishes to consume content/metadata from content or aggregation nodes in the network. For example, this could be an OPAC displaying reviews sourced via the platform or allowing customers to contribute reviews to an aggregation node.

6.2.3.3 Aggregation Nodes

Through BigFoot, described later, the Silkworm platform provides aggregation across multiple content nodes and application nodes. These are exposed in the same way as content nodes. In effect aggregation is a sharing mechanism that removes walls between nodes and presents them as one.

Aggregation nodes are a logical sharing, the physical sharing mechanisms can be diverse and mixed such as a physical union database or federated search or a hybrid. The key to aggregation nodes value is that the client making use of the node does not need to know or care how the physical sharing is occurring.

6.3 Silkworm Platform Architecture: Services

The following are key principles that drive the Silkworm architecture:

- Lower the technical and cost barriers to discovering, sharing and consuming content
- Centralise complexity into the platform components and hide it behind simple interfaces so removing complexity from application development
- Allow developers to work in the way they are most comfortable with, e.g. providing a range of protocols (e.g. REST, SRW, z39.50, UDDI) and language bindings (e.g. java, .net, PHP, perl)

-
- The network effect of problem solving e.g. if one person defines a useful transformation, others should be able to share it (if the originator wishes it)

6.3.1 Directory services

Directory services centralise the complexity of discovery and maintenance.

Since content nodes in the hidden web are not hyperlinked they cannot be discovered in the same way as web pages. Therefore the directory provides information to discover and access content. The core of the directory is institution, collection and service descriptions.

The directory incorporates revision control to allow roll back of individual entries. This allows for community style methods of entry maintenance to be combined with more controlled forms. This is the method used by 'wikipedia', an online encyclopaedia that anyone can edit.

The directory also incorporates notification mechanism via email or RSS to allow human and machine interfaces to remain aware of changes to entries they have an interest in. For example, an application may well cache directory entries locally for performance reasons.

The application can monitor the RSS feed for the cached entries and update the cache when a change occurs.

The directory is itself an aggregation node aggregating directory entries from multiple sources. As before, the value of the Silkworm platform increases according to how inclusive the directory becomes.

6.3.2 Access services

Access services centralise the complexity of multiple protocols by decoupling the consumer from the publisher.

To lower the technical barriers to discovery, sharing and consuming content it is vital that differences in technical access methods for each content node be hidden from the content consumer. The access service provides multiple consumer side protocols so that the consumer can choose the most appropriate one for their need (e.g. REST, web services, Z39.50). The purpose of the access service is to bridge whichever consumer side protocol is used to the required protocol for the content node.

For example, any Z39.50 content node will be accessible using simple SRU through the access services. Access services act as a proxy for the content consumer. Alternatively the consumer can contact the content node directly using the service descriptions in the directory.

6.3.3 Aggregation and sharing services – BigFoot

Aggregation services centralise the complexity of removing walls and sharing content across many nodes.

Bigfoot is the part of the Silkworm platform that provides aggregation services that can be used to create many forms of content sharing. A simple example would be moving Talis' current shared cataloguing system, Talis Base, onto the Silkworm platform. Essentially Talis Base is an aggregation over all the library catalogues of all the Talis library customers. Due to the fact that sharing involves contribution as well as consumption, strong network effects can be gained.

The Bigfoot aggregation services are a generalisation of the shared cataloguing model to other forms of content or metadata. This forms the basic infrastructure on which an industry wide architecture of participation could be built. Rich user created content such as community

information, reader reviews, librarian created reader development content, and student created reading lists (to name but a few) could be surfaced in all sorts of user applications.

6.3.4 Grouping services

Grouping services centralise the complexity of creating and working with views or subsets of all the nodes in a network.

It is often useful to draw circles around nodes to create groups that can be addressed as one or to create a particular view based on some criteria like geography, interest, subject or membership. Sometimes institutions maybe happy to share with certain other institutions but don't want information generally available. Grouping provides mechanisms for these scenarios.

6.3.5 Notification services

These centralise the complexity of maintaining synchronisation and currency across distributed loosely coupled systems both human and machine based.

Machine or humans can subscribe to notification of changes to directory content or aggregation content via methods such as RSS. This allows numerous maintenance and sharing scenarios to be created. For example, an application may want to keep a copy of service descriptions locally for performance reasons but need to update these if they are changed in the directory.

6.3.6 Transformation services

These centralise the complexity of transformations and allow sharing of useful transformations.

There are a number of different scenarios that may benefit from certain transformations of data. It may be that a certain content node search query needs transforming to optimise the results. To lower the barrier for useful systems built on Silkworm a general XLST transformation service is available to convert results into different formats or even into HTML to produce human interfaces. Transformation of search or harvest queries for particular nodes or node types can also be specified. It is vital that useful transformation can, if wished, be shared with other Silkworm users.

6.4 Silkworm's Platform Architecture: Participation

6.4.1 The need for openness and an architecture of participation

Put simply, the value of Silkworm is driven by the applications and the users of the platform and the community of developers that build those applications. It serves no one for the platform to be proprietary and closed.

In fact, Silkworm not only allows applications to be built with powerful architectures of participation, but the platform itself has a strong architecture of participation. This participation exists at many levels from creating new directory entries or adding new classification tags to existing ones (in the manner of Folksonomies) or developers contributing useful transformations or open source applications code.

6.4.1.1 How to share content: Creative Commons

Creative commons is one mechanism that can be used when contributors are happy to share content such as directory entries, classification tags or transformations.

6.4.1.2 How to share code: Open Source

Open source will play an important part in the platform and developer community. The Silkworm team will be seeking to work with leading open source companies to ensure as many developers can participate in the platform as possible.

6.4.1.3 How to share contracts: Open Standards

Nothing limits developer participation more quickly than closed proprietary standards. Silkworm will be built on open standards both existing and new.

6.4.2 Bigfoot aggregation services

The following example concerns book reviews but is applicable to any user-contributed content.

A single installation of an application has a limited number of users. For example, a library management system may allow reviews to be submitted but is limited to the reviews written by patrons of that library only. There are islands of small groups sharing reviews with walls between them. In fact the walls are so tight it may not be worth doing for many libraries because they do not have enough users and only a small percentage are likely to actually participate.

Bigfoot aggregation services can solve this problem at two levels.

6.4.2.1 Removing the walls between application instances

To increase the number of reviews and therefore the richness of the user experience, all the installations of that application running in the world could contribute reviews to a single pot. This is an aggregation over all the nodes running that particular library application and removes the walls between instances.

Bigfoot services on the Silkworm platform massively lower the technical and cost barriers for doing this. In effect this creates value directly from installed market share; the more installed applications the greater the value.

Ultimately this is still limited by the size of the install base for one application. Talis believe the real prize is the solution below.

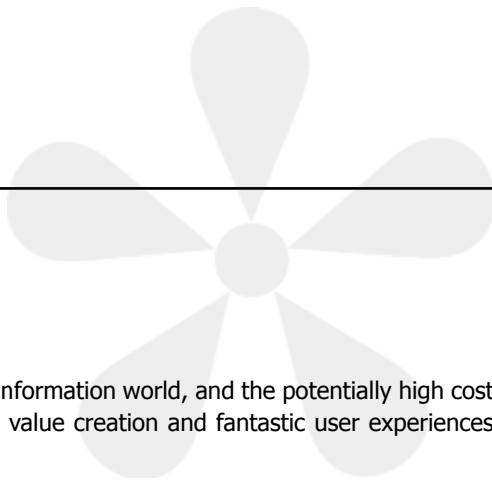
6.4.2.2 Removing the walls between different applications

What if any book related application could allow users to read and submit reviews into a single shared pot no matter whose application they were using? This would be truly all encompassing and without walls. This would enrich every application and have the widest possible participation.

6.4.2.3 Bigfoot – Participation without walls

Bigfoot aggregation services, built on open Web 2.0 standards, are easily integrated in any modern application (web based or thick client) to allow many different sharing and aggregation scenarios.

Bigfoot can be used for all encompassing aggregation or, in combination with the Silkworm grouping services, it can be used for sharing and aggregation within specific boundaries such as only across a vendor's own applications, specific interest groups or between members of a consortia group.



Conclusion

The multiple silos of content in the library and information world, and the potentially high cost of exposing them, is preventing the creativity, value creation and fantastic user experiences that should be possible today.

The limitations of the situation are increasingly worrying as it becomes clear that the value that can be unlocked by breaking down the barriers between islands of information and making content work in many new and different ways is enormous.

Technically, there's no reason for this situation to continue. The age of walled gardens and proprietary closed systems is passing and the universal connectivity and massive power of today's machines mean that everything discussed in this paper is already happening. It's a matter of how and when, rather than if.

Perhaps the greatest barrier to progress, then, is the paradigm shift that must take place in the library and information systems world. Instead of traditional isolationist approaches, it's time for vendors to collaborate as partners rather than competitors.

It's a fundamental change in our industry but one that will pay great dividends. If you or your company would like to be involved or wish to find out more then please visit silkworm.talis.com. Here you can join in the discussion via the Silkworm discussion forums and find presentations, demonstrations and many other ways to become involved.

About Talis

Talis is the leading provider of library management solutions for the UK and Ireland. Designed for the academic and public sectors, Talis' suite of products helps public authorities, universities and colleges to manage their extensive resources and provides users with swift and easy access to that information. Thanks to Talis' support for greater integration, libraries in both the public and academic sectors are also able to streamline multiple systems and improve internal efficiency.

Founded in 1969, Talis has grown steadily and now enjoys an impressive 25% share of the UK library market. The company also has a long history as a technology innovator and invests heavily in research into the information management technologies of the future. Ongoing research projects include metadata, web services, RSS and service-oriented architectures.

For more information, please visit www.talis.com

About the Author

Justin Leavesley - Chief Technology Officer

Justin joined Talis in January 2000 from Morse Group plc, a leading European technology supplier. After receiving his Master of Physics from Oxford University, Justin joined IBM working first in the defence and aerospace sector for the British Army and later for Natwest Bank.

Justin left IBM to join an internet start up company based at Birmingham University where he was involved in product development, marketing and business strategy. Justin, a Sun certified Java enterprise architect, is now responsible for technology and related business strategy at Talis where he leads the Research Group.