

talis



Supporting the next generation of applications for delivering rich, library content and services



A white paper by
Ian Davis

Talis

[Version 1.0] November, 2005

Contents

Executive summary	3
1. Web 2.0 and the Amazoogole Effect	4
2. What are next generation applications?	5
The history of the platform	5
3. What does this mean for the future of library systems	6
Directory Services – Project Silkworm	8
Semantic Content Management – Project Bigfoot	9
Integration Services - Projects Keystone and Arpaz	10
Transaction Gateway – Project Stargate	11
4. How the Talis Platform could transform the library domain	12
Cataloguing and holdings	12
Inter Library Loan services and sourcing	12
The value of aggregated data	12
Conclusion	13

Executive summary

Why can't using a library catalogue be more like the experience you receive when you visit Amazon or Google? User-centric, personalised, easy to use, these sites are at the heart of the high expectations that many people now hold about finding information. So why can't the same rich, fulfilling experiences be available to users of applications that provide services to institutions?

The answer is simple. There is no reason why this should not be the case.

Sites such as Amazon and eBay are based on the idea of Web 2.0, a concept of delivering rich, participative, personalised experiences through open and economically viable applications. These next generation applications can be applied to the library domain just as they can to any other.

This is exactly what the Talis Platform supports. It is a revolutionary technology that will have a dramatic impact on the way information is shared and used by individuals and institutions. It works by providing the underlying infrastructure to support the collection of a large pool of content that is not held in the usual information silos but freely available to end users and other content providers. Libraries will be able to draw on this mass of information and deliver it to their users more cost effectively than ever before.

1. Web 2.0 and the Amazoogole Effect

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 can only get tougher as these companies expand their reach into more aspects of people's lives. If libraries are to thrive, as we all believe they must, then their systems must also deliver on these expectations.

There is currently a lot of hype and discussion around the term 'Web 2.0'. Several people have attempted to define Web 2.0 (and the previous Web 1.0), but looking at some of the latest (usually consumer oriented) sites on the Web, we can discern a number of clear trends within this concept:

- **Participation** – sites such as Amazon, where individuals can contribute book reviews, and eBay, where purchasers rate the sellers (and vice versa), have demonstrated the huge value that is generated when many people all participate and contribute to a single site.
- **Open** – this term can be applied both technically and economically. It is technically very easy to join a particular community (as a contributor, consumer, or both), and usually at minimal or no cost.
- **Rich** – a modern web site contains material that can be described as 'rich'. Using Amazon as an example again, not only does the user get the title, price, and availability of a particular book (or CD, or DVD etc.), but they also see an image of the product, consumer and publisher reviews, extracts (sample chapters, 30-second excerpts), and a lot more.
- **Fulfilling** – the goal of many modern web sites / services is to provide a fulfilling experience for the user. For example, if a user is looking for information, Google provides the search mechanism and a direct link to the original document (and even a local cached copy in case the original document has disappeared between the time that Google indexed it and the time of your search). eBay, not only facilitates the auction transaction, but through its PayPal service, the associated payment transaction as well.
- **Personalised** – Web 2.0 sites learn from an individual's interactions and then use that learning to provide personalised content – for example recommending particular items based on what you have already purchased.
- **Remixing** – another trend that is emerging is the remixing of data and applications to create completely new applications and additional value. Often this remixing is without the involvement of the original applications and data owners. There are for example now several sites that mix data with Google Maps to create a geographical view of data.

What we are describing here are the characteristics of the next generation of web-based applications. Currently, most of these characteristics are predominantly seen in consumer-oriented sites. But, there's no reason why the same experiences should not be available to users of applications providing services to institutions.


2. What are next generation applications?

The history of the platform

Within the technology world, the term platform is used to describe hardware and / or software on which applications are built. For example Wintel (Windows + Intel) is the platform familiar to all of us on the PC.

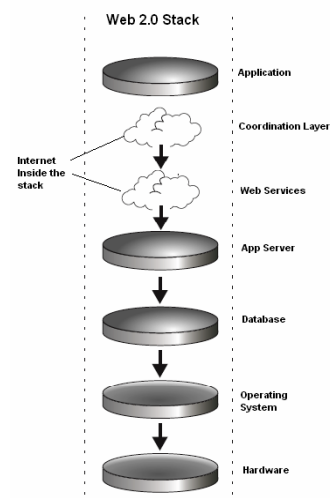
In general a platform is there to hide the complexity of the layers below, and to provide a set of common services. The traditional computing architecture has consisted of a series of platforms built on top of each other:

- In the early days of computing the platform was the hardware itself, and applications were written directly in machine code – the language understood directly by the hardware.
- Fairly soon, the concept of an operating system was added. The purpose of the operating system was to let the programmer concentrate on building business functionality and not worry about the particular hardware. So, for example, whether a file was stored on a magnetic disc, a tape, or a CD-ROM should be transparent to the program.
- With the advent of large-scale data processing, a third layer was added to the stack – the Database.
- During the 1990s a fourth layer was added – the Application Server that provided a richer container for building applications.



The next generation of applications

This fundamental picture did not change even with the emergence of the Internet and the World Wide Web. Applications were still written on the four-layer stack, but may have been accessed via the Web, or communicated with other applications via the Internet



However, in the latest generation of applications, this has changed. Now, applications are built on top of the Internet. In other words, the Internet is inside the stack. In this architecture, applications are built by making calls to a series of Web Services, which may be running, on any machine, anywhere. The most important implication of this is that to build (and run) applications, you no longer need to buy and install hardware, operating systems, databases, and applications servers. An application can be built and run in a web browser running on a simple PC connected to the Internet. Universal connectivity, open standards, and rapidly improving price/performance are rapidly changing the way we view applications.

The 'network effect'

The other key characteristic of the latest generation of platforms is that they provide not just services, but also rich, shared data. Therefore as more and more applications use that platform, the value of the platform itself is increased through what is known as the network effect.

3. What does this mean for the future of library systems?

Talis has spent a considerable amount of time investigating ways that libraries could use next generation applications to improve the experiences they offer users. Our conclusion is that to reap the benefit of these technological changes, there must be several fundamental shifts in the way that information technology supports libraries.

Collaboration

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.

Sharing

A shared platform is required to lower costs, break down walls and provide an all encompassing experience for users. The market today 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. It does not make economic sense for each player to try to provide the all encompassing experience enabled by next generation applications, and the market should not have to bear the weight of each vendor duplicating this.

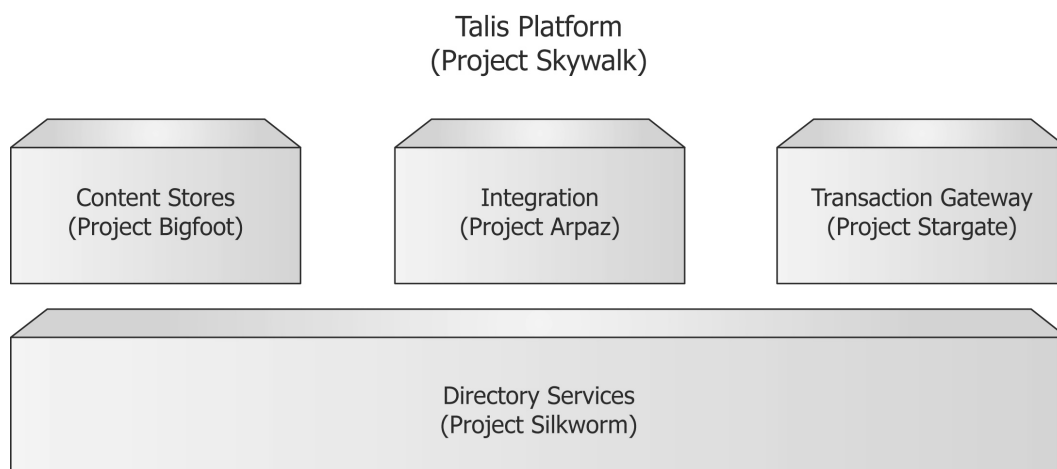
The Talis Platform is an answer to these challenges and will bring about a fundamental change in the cost structure of the industry that will enable libraries to thrive and to deliver exceptional value to their patrons.

4. What is the Talis Platform?

The **Talis Platform** is a revolutionary new architecture for delivering applications that provide high value at low cost, using the latest technology of the World Wide Web.

The Talis Platform is an open, extensible, shared, 'Web 2.0' platform that initially consists of four main components:

- Directory Services
- Content Stores
- Integration
- Transaction Gateway



Unlike many traditional platforms, there is nothing to install when using the Talis Platform. The platform is accessed by applications using a series of Web Services. The platform itself is generic and over time there will be a number of instances of the platform, each focussed on a particular information domain, but linked at the underlying platform level.

The first such instance is one focussed on the Library domain. This is no surprise when you consider the extensive knowledge and experience that Talis has in this field.

The value of any platform is increased by the existence of an ecosystem of developers and other partners using that platform to build new applications, or to enhance existing applications. Through the Talis Connexions program and Talis Developer Network, partners are being recruited to this network. Some are traditionally competitors of Talis in the delivery of Integrated Library Systems, but can see the benefits to all involved by using a common underlying platform.

Directory Services – Project Silkworm

At the heart of the Talis Platform is a global, open directory and set of associated services. The Talis Directory is in effect the map that describes the location and access methods of all the resources known to the platform – the content stores, contributing nodes, applications nodes and so on. One characteristic that makes the Talis Directory distinctive is that it is 'self maintaining'. In other words, there is no central body responsible for maintaining the directory – instead responsibility is distributed to those who are best placed to perform that maintenance. Built in mechanisms continually check the validity of data and send alerts if problems are found.

In the library domain, for example, the platform directory contains an entry for every known collection, including the URL and other details of its OPAC. The address (URI) of the directory entry for that collection is guaranteed never to change. However, individual libraries, for any number of very valid reasons, may have to change the details of their OPAC. By updating just one directory entry, in the Talis Directory, any application, anywhere, that accesses the OPAC via the Talis Directory, will continue working when the details change. On the other hand, if an application has hard coded the OPAC details, it will cease to work until manually changed.

For libraries, the huge benefit of the Talis directory is that there will be just one place where they need to maintain their library details such as address, opening hours, website, and OPAC. Then any application on the web that needs such details will be able to access them through standard Web Services.

Semantic Content Management – Project Bigfoot

Historically, data has been held in databases and accessible only to those applications specifically designed to access it. This has meant that such data has not easily been shared and whatever the potential value of that data, its actual value is limited by the capabilities of these specialist applications.

The Talis Semantic Content Management subsystem has been designed to overcome this very significant restriction and to make content as widely available as possible. The key characteristics of this part of the Platform architecture are:

- **Discoverable** – Bigfoot makes use of the Directory Services described above to publicise itself and ensure that any application needing access can find and use the content (subject, of course, to appropriate security and authentication).
- **Searchable** – the content of the store is searchable through general purpose, open interfaces. This means that applications do not need to know anything about the internal representation or organisation of the data in order to be able to locate it. The platform includes built in indexing and searching mechanisms that are independent of the content. This means that even if the content is not hosted on the platform, it can still be discovered efficiently.
- **Shared** – one of the features of the new Web 2.0 world is the concept of shared data. Users and applications do not need to keep their own, local, copy of data, but access a common, shared version as and when they need it.
- **Scaleable** – Bigfoot has been designed with scalability in mind. There are no built-in limits to the amount of data that can be held, and a single store can be physically distributed across many machines to ensure that performance is not affected by size.
- **Flexible** – the underlying design makes no assumptions about the nature of the data that can be stored. Therefore whatever the data, structured or unstructured, text, images, sounds, videos, maps and so on, it can be stored, searched, and retrieved efficiently.
- **Economical** – perhaps the most dramatic characteristic of this architecture is that it has been designed to minimise the costs of providing, storing, and consuming data. Smart use of the latest technology, including low cost commodity hardware and bandwidth, and open source software, minimises the cost for all concerned.

All of the above enable individual providers or consumers of data to focus on their core competencies and do not have to worry about where or how to store data. The Bigfoot mechanism in effect 'decouples' the various parties.

Using this technology, the Talis Platform will enable applications to provide rich user experiences. Resources will be easily discoverable and presented with enrichments and links to other resources.

Integration Services – Projects Keystone and Arpaz

The integration capabilities of the Talis platform are the key to making it accessible to all. For many years Talis has been heavily involved in the world of standards, and is a member of many of the world's most important standards bodies such as the World Wide Web Consortium (W3C), the Dublin Core Initiative, NISO, VIEWS etc. Wherever possible, Talis uses existing recognised open standards. In other cases, Talis will propose new standards and work with the appropriate bodies to get them adopted.

Whether contributing or consuming content and services, the aim of the Talis Integration Architecture is to make the job as simple as possible in order to minimise the friction of doing business.

As an example, any node wishing to contribute data to the platform need only install a small, low cost, integration client. This flexible piece of software can either feed data from the contributing node to Bigfoot, or if preferred, make the content of the node accessible to other applications on the network. In this way, the owner of the contributing data retains full control over whether that data is copied or referenced by the Talis Platform.

The Talis Integration Architecture is also designed to make it much easier for partners, such as applications vendors, to 'mix and match' their software to deliver increased value to their customers. It is a step away from the idea of large, monolithic, unwieldy systems.

At the individual library node, the Integration technology will enable full integration of the integrated library system and platform services with institutional workflows. In a public library environment this could involve integration with the local authority portal whilst in the academic environment it could be integration with the student registry system.

Transaction Gateway – Project Stargate

The Transaction Gateway component of the Talis Platform is the underlying mechanism for enabling transactions between business partners. It is being built based on the extensive experience that Talis has with its Talis Gateway product for libraries. This product provides EDI based services for transactions such as quotes, orders and invoices between libraries and their suppliers. The new Transaction Gateway will continue to provide all of the existing functionality, but also new transport mechanisms (specifically the Internet/Web) and transaction types.

As with all other components of the Talis Platform, the Transaction Gateway will be standards based and open. It will be technically easy and economically viable to connect to and use.

Being part of the Talis Platform, the Gateway will also have access to content that can be used to enrich the transactions, and via the directory, details of participating institutions. The need to maintain a lot of common data separately will therefore be minimised.

Within the library domain, one of the more exciting aspects of the Gateway will be the automatic integration of the procurement processes into other library functions. For example library holdings information on new stock will be fed into Inter-lending agencies. Stock performance information will be generated and used to support the selection process. National or regional reading trends from Amazon will be used to further inform the selection process.

The Gateway will enable richer selection, and more choice when purchasing outside of the main supply contracts, using price comparisons for cheapest supply facilities. It will support single, regional as well as national consortia regardless of ILS vendor.

The Gateway is all about the delivery of transactions and orchestrating the simultaneous activities that support those transactions. Full audit, monitoring and chasing capabilities will be built in giving guaranteed confidence and information at the required times.

4. How the Talis Platform could transform the library domain

The Talis Platform will have a significant impact on the library domain.

Cataloguing and holdings

For many years, Talis has provided cataloguing and holdings capabilities to users of its ILS and other libraries. These services (Talis Base and UnityWeb) have now been upgraded to utilise the Talis Platform, providing considerably more power and flexibility than was previously possible.

- Talis Base 2.0 has been uncoupled from other Talis applications and is now available to any library, whether as a contributor or as a consumer of data.
- The number of datasets provided has been dramatically increased and now also covers other non-book formats such as CDs and DVDs.
- The new technology has enabled Talis to dramatically improve the importance of its holdings.
- Data enrichments such as book jackets, tables of contents, and reviews are now included.

The most dramatic application of the platform will be the provision of a free web-based resource discovery service. It has long been argued that libraries should not have to pay to access their bibliographic and holdings data that they have contributed to various union catalogue initiatives. Well now, the economics of the Talis Platform will enable such a service, free to anyone, including libraries, with access to the Internet. This will greatly help libraries to expose their valuable collections to the world at large.

Inter Library Loan services and sourcing

An inter library loan (ILL) request management system is being implemented on top of the Talis Platform. Thanks to the platform economics and open standards being used, this will be available at a much lower cost than traditional systems. And as more and more institutions join the platform, the network effect will dramatically increase the value of the solution. The Talis Platform will also enable alternative methods of sourcing items for a borrower. For example, through tight links to web sites such as Amazon and eBay, books or music could be bought rather than sourced from another library – possibly at a much lower cost.

The value of aggregated data

Possibly the biggest value of a data-centric platform such as this is the information contained within the aggregation of data from many participants. The Platform will support advanced data analysis and data mining capabilities to support the following:

- Generation of aggregated statistics to support regional and national reporting for government and other bodies
- Support for aggregated stock management by revealing borrowing trends across any slice of the market
- Optimisation and automation of the library book supply chain via integration with the Transaction Gateway component.

Conclusion - Richer library user experiences

The Talis Platform is a revolutionary new technology, available now, that will have a dramatic impact on the way information is shared and used by individuals and institutions.

Using the latest technologies and principles from the emerging Web 2.0, the Talis Platform provides a very open, economic and data-centric base on which to build next generation applications.

Many components of the initial Talis Platform, which has been designed for the library domain instance, are already delivering value to customers. In the future, all of the Talis applications will be upgraded to take advantage of the Platform features, providing a richer user experience at a more economical price.

If you would like to receive more information about the Talis Platform or you would like to participate, please contact Dave Barker at dave.barker@talis.com or call: 0870 400 5000.

The document is intended for the general reader who is interested in what will be possible with the next generation of applications in the library domain. For a more technical view, please refer to the forthcoming technical papers on each component.



Copyright 2005 © Talis

www.talis.com