
Games of tag: The rise and disappearance of enterprise collaborative bookmarking

Matt Moore*

Social bookmarking sites such as Delicious have arisen on the internet. Enterprises have been experimenting with open source and commercial versions of such tools within – and sometimes across – their firewalls. This article begins with an overview of social bookmarking and folksonomies. It then explores two case studies (MITRE Corporation in the United States and New South Wales Department of Education and Training here in Australia) and two vendor offerings (Connectbeam and IBM Lotus Connections). It ends by offering some conclusions on the future of enterprise collaborative bookmarking and some advice for those wishing to engage in it.

INTRODUCTION

Information management tools are developed because users have specific needs in specific contexts. If those contexts and needs change then the development of new tools may be required. For example, the spread of SMS messaging has led to the creation of new (and not necessarily attractive) forms of language and expression to allow quick communication. To make things even more complicated, tools developed for one purpose may have unforeseen applications – eg who would have predicted that something as simple as text messaging would combine with reality TV to create an interactive experience?

In the realm of metadata (although it was not called that then), librarians developed card catalogues to allow readers to find books in libraries without having to visit every shelf and classification schemes to allow effective browsing. Unintended consequences occurred. Unlike the Dewey Decimal System, the Library of Congress Classification Scheme was not created to map the whole of human knowledge; its letters and numbers refer to specific locations in one particular library. Only as the prestige of that library increased did it take on a global role.

Contexts change and the last decade has seen a proliferation of electronic information artefacts both inside organisations and in the public sphere, including books, articles, spreadsheets, presentations, blog posts, tweets, wiki pages, images, podcasts and videos – the list goes on. How do we manage this torrent of stuff? On the internet, Google's PageRank algorithm provides one way of combining human behaviour and machine processing power to improve the findability of information. It is far from perfect but until those who advocate the Semantic Web can deliver on their promises, it remains the best option.

However, the success of Google presents organisations with a problem. Employees and managers see the power of Google on the internet and say "I want that". So IT managers take this request literally and purchase Google Search Appliance. The results are often unsatisfying. There are several reasons for this but a key one is that Google's power comes from the highly-hyperlinked nature of the internet. The PageRank algorithm depends upon aggregating the judgments of others. In stark contrast, organisational information environments are often highly unhyperlinked. Information artefacts stand in splendid isolation from each other. Those that do contain hyperlinks often do not work because the files they refer to have been moved and thus the links broken.

So what is an information manager to do? Part of the solution may come from another internet-based information management tool. For many years now, web browsers have allowed users

* Matt Moore is a director of Innotecture, an occasional lecturer at the University of Technology, Sydney, and chair of the New South Wales Knowledge Management Forum. He has spent over a decade working in knowledge and information management, learning and development, and internal communications with organisations such as PwC, IBM, Oracle and the Australian Securities and Investment Commission.

to bookmark webpages and to manage those bookmarks in folder hierarchies. This is useful for the individual user but does not scale up. How do I share these bookmarks with other people? And what do I do if a bookmark belongs in more than one folder?

In 2003 a service called *del.icio.us* was launched (Delicious is now owned by Yahoo). Delicious was a website and a plug-in for your browser. You create an account and when you find a site you may want to visit again, you can add it to your list of Delicious bookmarks at the click of a button. If you visit the Delicious site you can view all of your bookmarks; and unless you have marked that bookmark “private” then so can everyone else. When you create a bookmark, you are also prompted to add some short descriptive text in the “Notes” field and some tags in the “Tags” field. Tags are freetext keywords. Delicious does not offer a controlled vocabulary; all tags are defined by users. Delicious only supports tags that consist of individual words rather than phrases – ie it would interpret “car park” as “car” and “park”, therefore Delicious users would use the tags “carpark”, “car.park” or “car_park” instead. Other social bookmarking tools do support phrases as well as words.

FIGURE 1 The Delicious prompt box

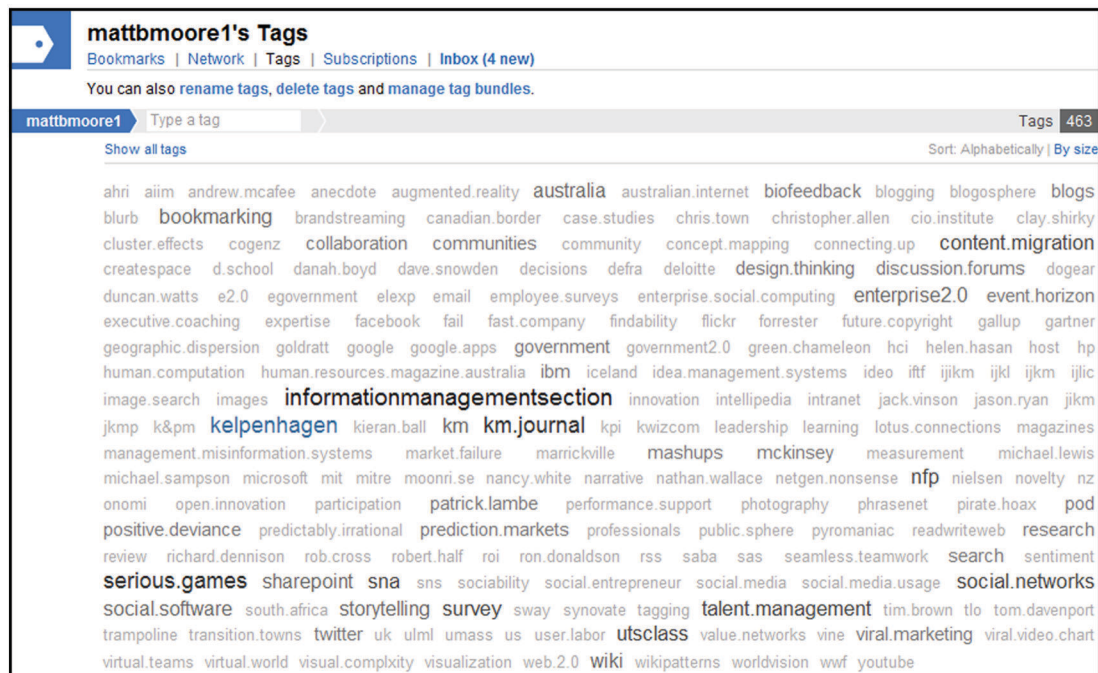
The screenshot shows the 'Save a Bookmark' dialog box in a browser window. The dialog has a title bar 'Save a Bookmark' and a close button. Inside, the 'delicious' logo is in the top left, and 'Signed in as mattbmoore1' is in the top right. The main area contains several input fields: 'URL' with 'http://www.smh.com.au/', 'TITLE' with 'Sydney Morning Herald - Business & World News Australia | smh.com.au', 'NOTES' with a '1000 chars' limit, and 'TAGS' with a hint 'Separate Tags with spaces. E.g., hotels bargains newyork (not new york)'. There is a 'Mark as Private' checkbox. Below the 'TAGS' field is a 'SEND' button with the text 'Send your bookmarks' and a 'Clear' button. A 'Tags' button is also present. Below the 'Tags' button, there are two sections: 'Recommended Tags: click to add from your existing tags' with tags 'australia' and 'blog', and 'Popular Tags: click to add from popular tags on Delicious' with tags 'news', 'australia', 'newspaper', 'sydney', 'newspapers', 'media', 'daily', and 'smh'. At the bottom right are 'Save' and 'Cancel' buttons.

In addition, Delicious offers two prompts:

- The “Tags” prompt shows two kinds of tags:
 - “Popular Tags” that others have applied to this URL; and
 - “Recommended Tags” that the user has used with other URLs before.
- The “Send” prompt allows users to inform others of this bookmark. These can be sent via emails or broadcast via Twitter to Delicious members that the user has marked as being in their network.

Users can also receive updates when a specific tag is used by others, when a specific Delicious user in their network adds to their bookmark collection or some combination of the two.

Using Delicious also has an unintended consequence. The frequency with which certain terms are used by individuals or groups can be measured and represented numerically or visually. A tag cloud is such a visual representation, with the size of a tag being related to its frequency. The author’s personal tag cloud is included in Fig 2.

FIGURE 2 The author's personal Delicious cloud

As shown in Fig 2, certain tags are more common than others. Some care needs to be taken in interpretation though. The most common tag is rather obscure. In fact, here it was used to mark social software resources for a consulting client – it seemed more appropriate and dynamic than copying URLs into a Word document. While that tag has meaning to the interested parties it is not useful for others.

At a personal level, a tag cloud can indicate a person's interests. At a collective level, it can offer an insight into shared issues and culture. At a collective level, a tag cloud can be seen as folksonomy, a term coined by Thomas Vander Wal to describe the emergent metadata structures that proliferate on the internet. Folksonomies do not have the rigour of controlled vocabularies, thesauri or taxonomies but they are cheaper to create and maintain and can also give insight into the subjective experiences of users/creators. In dynamic, ambiguous environments, folksonomies may offer an advantage over more structured or controlled alternatives. However, as each offers such complementary opportunities, information professionals should learn to work with both.

So what does this have to do with the world of work? Several organisations have been experimenting with collaborative bookmarking and consequently a range of open source and paid-for products has emerged to meet these needs.

More information is available at:

- Delicious (<http://www.delicious.com>);
- Folksonomy (<http://www.vanderwal.net/folksonomy.html>);
- Golder S and Huberman BA, "Usage Patterns of Collaborative Tagging Systems" (2006) 32(2) *Journal of Information Science* 198 (<http://www.hpl.hp.com/research/idl/papers/tags>); and
- Chapter 10 of *Organising Knowledge* by Patrick Lambe (Chandos Publishing, 2007) (<http://www.organisingknowledge.com>).

CASE STUDY: THE MITRE CORPORATION AND ONOMI

The best documented case study of collaborative bookmarking in organisations belongs to the American-based MITRE Corporation. MITRE is a not-for-profit research organisation that employs over 7,000 scientists, engineers and support specialists. In 2005, a group of researchers and corporate

services staff began investigating the utility of collaborative bookmarking within their corporate environment as a way of sharing information and building communities of practice. They selected an open source tool called SCUTTLE because of its ease of installation, extensibility and existing features set. They used SCUTTLE to build a tool known with MITRE as Onomi. This tool was developed incrementally and iteratively with user feedback incorporated into each design increment. This led to the addition of new features such as the ability to email bookmarks, display related users, browse by user, interface with other programs and search by corporate organisation affiliation (via MITRE's LDAP server) and by file type.

One of the issues with social software is that users tend to generate most of the content. However, users are more likely to contribute content if there is some already present. This puts social software designers in a chicken-or-egg situation. The Onomi team used a number of strategies here. They approached influential information-sharing employees with large internal social networks (such as librarians). For librarians, Onomi offered an immediate benefit: rather than send a collection of resources on the same topic to multiple users (eg phylogenetics), librarians could simply send them the relevant tag list. Onomi also allowed users to import their Delicious bookmarks into the system – creating an immediate source of metadata. Onomi allowed simultaneous posting of tags to itself and Delicious (which some internal URLs blocked from external posting) to make life simple for corporate users who wanted to maintain a bookmarking presence both inside and outside the organisation. In addition, the team promoted Onomi through a range of corporate communications channels, including intranet banner ads, demonstrations, briefings and giveaways of real bookmarks.

Some statistics provided by MITRE as of September 2008 were:

- almost 2,000 employees had visited Onomi;
- 275 users provided content;
- 12,600 resources were bookmarked;
- over 9,700 unique tags were added 68,000 times; and
- as of 2006, around one-fifth of all resources bookmarked were internal to the enterprise.

So what were MITRE staff using Onomi for?

- Personal bookmarking – users were more likely to examine their own bookmark collection (70%) than those of others.
- Information sharing and dissemination – one project manager told his project members that all papers, works-in-progress, deliverables and resources relevant to the project should be given a project-specific tag. He stated that “the value of Onomi is disproportionate to my amount of use”.
- Information discovery – when looking at bookmarks by topic, users were more likely to look at all bookmarks by tag (75%) than their own (22%).
- Forming communities and social networks – users stated that they changed their tags or added new tags based on the tags that others with similar interests were using.
- Expertise finding – the Onomi team noticed that one user had been rapidly bookmarking select employee directory pages with the term “expert” as a tag along with a skill or research area. Another user must have noticed this because he/she bookmarked his/her own directory page with the terms “expert” and “design”.

More information is available at:

- Damianos L, Griffin J and Cuomo D, *Onomi: Social Bookmarking on a Corporate Intranet* (Mitre, 2006) (http://www.mitre.org/work/tech_papers/tech_papers_06/06_0352/06_0352.pdf);
- Damianos L, Cuomo D, Griffin J, Hirst DM and Smallwood J, *Exploring the Adoption, Utility, and Social Influences of Social Bookmarking in a Corporate Environment* (40th Annual Hawaii International Conference on System Sciences, 2007) (http://www.mitre.org/work/tech_papers/tech_papers_06/06_1199/06_1199.pdf); and
- Cuomo D, Damianos L and Griffin J, *Exploring Social Bookmarking* (Mitre, 2008) (http://www.mitre.org/news/the_edge/summer_08/cuomo_damianos_griffith.html).

VENDOR PROFILE: CONNECTBEAM

MITRE built their collaborative bookmarking tool themselves using open source software. However, there are a number of commercial vendors offering enterprise tools. Connectbeam is an American-based software startup focused on collaborative enterprise bookmarking. They began as a Delicious “behind the firewall” and, according to their CEO Sanjeev Sisodiya, did not fully understand what the uses for their product would be.

Their first customer was the engineering division of global manufacturing company Honeywell. They began with a small group of around 50 people. As with the MITRE example, they loaded the software with users’ bookmarks from Internet Explorer and Firefox to ensure there was content in the system from the beginning. When they investigated usage patterns they found that staff were primarily using the tool to search for people rather than content. Without intending to, they had built an expertise location system.

This encouraged them to make user profiles a key point of the application. Profiles could show the URLs bookmarked by users and the tags used to describe them. In addition, users could join groups based around specific topics. Users could tag content but they could also tag themselves with their expertise (the product does not currently allow users to tag others).

The critical question that Connectbeam and Honeywell had to deal with was how to embed the tool into the ordinary staff workflow. Connectbeam could be weaved into corporate search functionality via a browser plug-in and integration with the corporate search tool. As only 40% of staff used the internal search function it was also necessary to cross the firewall and support external searches as well. In addition, Connectbeam could offer social context around the content found – eg who else was interested in the same topics that the staff member was searching for.

As part of its product development, Connectbeam has built additional connectors to other information management tools. A connector is available for Outlook to integrate collaborative bookmarking with email. Connectbeam offers a “social business card”, an icon that allows the user to see the most recent bookmarks and tags of those they are communicating with. In addition, there is an application programming interface (API) to pull in user activities from other enterprise applications. This allows for richer user profiles and for greater relevance in the expert and document recommendations that Connectbeam provides. At the same time this can provide richer cues about their needs to others. Obviously it is also important to be able switch this functionality off.

Currently, Connectbeam offers connectors to Outlook, Sharepoint, Confluence and Google OneBox. Apart from Honeywell, Connectbeam has several clients in the United States pharmaceutical sector and is working for the United States Food and Drug Administration.

In the future, Connectbeam wants to offer a recommendation service for enterprise content similar to Amazon – “if you found this financial report useful then you should look at this one too” – based on the previous behaviour of other staff members. They also want to explore Connectbeam as a social network analysis tool.

For more information, visit: <http://www.connectbeam.com>.

CASE STUDY: NEW SOUTH WALES DEPARTMENT OF EDUCATION AND TRAINING, AND CURLS

The New South Wales Department of Education and Training (NSW DET) is responsible for all public schooling and public sector adult education within the state of New South Wales. The Centre for Learning Innovation (CLI) was established in 2004 to foster innovation in teaching and learning across NSW DET. In late 2008, CLI began piloting a tool called CURLS. CURLS is a collaborative bookmarking tool that could be considered a hybrid of a public tool such as Delicious and private tools such as Onomi or Connectbeam. Many of the bookmarks and tags entered into CURLS are aimed at the teaching and learning community and are publicly accessible on the web. However, bookmarks and tags can only be entered by NSW DET staff and teachers. Like Onomi, CURLS was built using an open source tool called SCUTTLE as a platform for development.

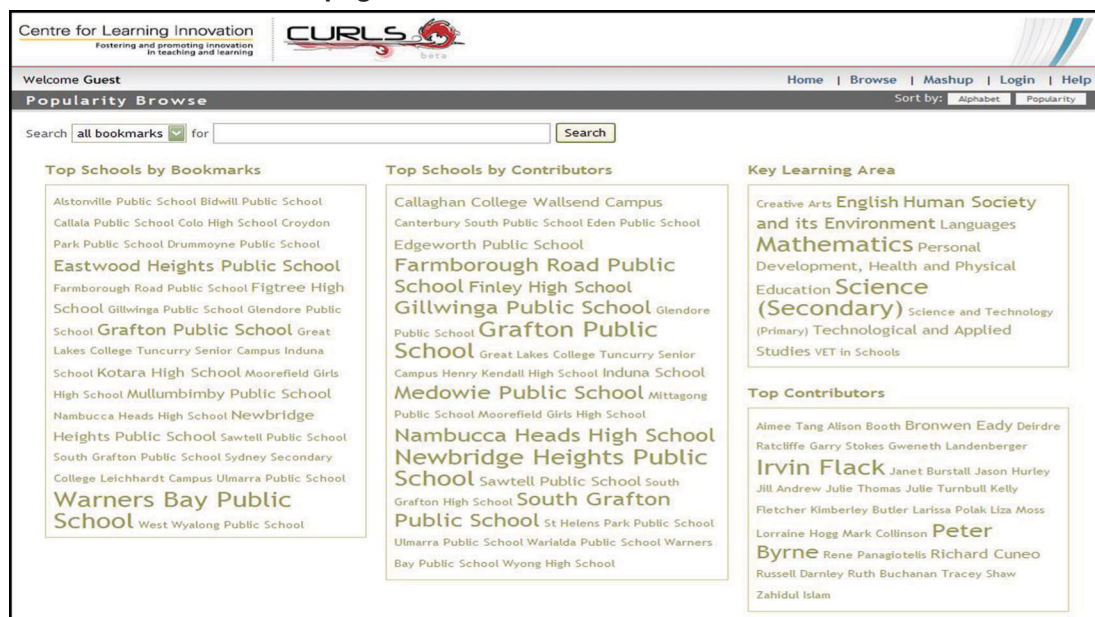
Dheeraj Chowdhury, Assistant Director for Learning Applications and Multimedia at CLI, explained that the impetus behind the development of CURLS was a need for educators to share online links with each other in a simple manner. Rather than creating a detailed taxonomy, CURLS allowed staff to use their own tags. One teacher is quoted as saying, “[CURLS] is fundamentally changing how I manage my online links, and [gives me] the ability to share these with other teachers and students”.

Some statistics provided by CLI as of December 2009:

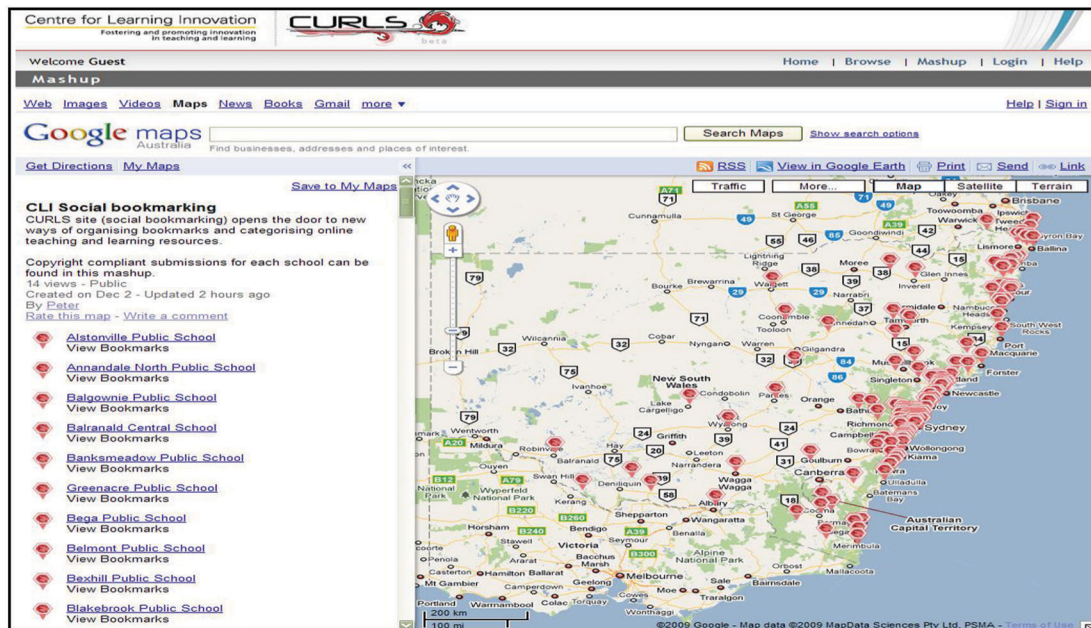
- over 15,200 teaching and learning links with 1,095 tags used;
- over 6,000 unique registered DET users from across 1,500 schools;
- over 700 contributing users (including 311 contributing users from 209 schools);
- the CURLS site receives over 1,600 visits per month; and
- around half of all bookmarks relate to the internet and half are NSW DET-based.

Since the initial rollout, CLI has been adding additional functionality to the tool. CURLS is linked to the NSW DET staff portal and so integrates with the portal’s single sign-on and filtering capability. Users can browse the tool by the most common tags, top individual contributors, top contributing schools and also by key learning area if added as a tag.

FIGURE 3 The browser page of CURLS



Users also have the ability to rate bookmarks on a scale of one to four (similar to Amazon’s rating system) and to view the most popular bookmarks in terms of ratings or hits. Recently a mash-up of CURLS with Google Maps has been made available. This allows users to see the physical locations of the schools that have contributed bookmarks and then view those bookmarks.

FIGURE 4 The mashup of CURLS and Google Maps

Many of the visitors to CURLS come from outside New South Wales. The CLI team wishes to extend the service for contributing teaching and learning links to education departments in other States, although at present there are no plans to open up bookmarking to those outside the education system.

More information is available at: <http://www.curls.edu.au>.

VENDOR PROFILE: IBM LOTUS CONNECTIONS BOOKMARKS

IBM's research labs developed a tool called "Dogear" that essentially provided the functionality of Delicious within IBM's internal information environment. As is common among IT companies, IBM's internal information environment is complex and diverse. The author's experience as an IBM employee was that while there was a lot of useful information to be found, locating it in a timely manner was another matter. Dogear provided a way of navigating this information jungle. IBM's Lotus software group was one of the first enterprise content software providers to market with a "social software" offering called "Lotus Connections".

Connections includes: functionality around profiles, communities, blogs, files, wikis and bookmarks. Connections also integrates with other IBM products such as Notes, Sametime, Quickr and Websphere Portal as well as Microsoft Sharepoint, Office and Outlook.

Connections Bookmarks offers the collaborative bookmarking functionality that you would expect plus it allows users to import their bookmarks from their browsers and to create watchlists of bookmarks created by a specific person or using a specific tag.

More information is available at:

- Millen D, Feinberg J and Kerr B, *Social Bookmarking in the Enterprise* (ACM Queue, December 2005) (<http://www.queue.acm.org/detail.cfm?id=1105676>); and
- Lotus Connections (<http://www-01.ibm.com/software/lotus/products/connections>).

CONCLUSIONS

Enterprise collaborative bookmarking can offer benefits beyond personal information retrieval that is independent of a single computer or browser – ie the WIIFM (What's In It For Me) that can initially attract users – but bookmarking also offers opportunities for improved team collaboration, expertise location and collective sense making. Bookmarks are corporate, semantic and social glue.

Bookmarking behaviours are a minority practice – but that is not necessarily a problem. A subset of employees will use the system (and this subset may increase over time). An even smaller subset will bookmark. There is some evidence that bookmarking behaviour follows a power law (ie a small minority or “short head” of bookmarkers produce the majority of bookmarks with a long tail of occasional taggers). However, a relevant article by Huberman and Golder¹ indicates that a small number of people can be representative of a broader grouping. They found that in public bookmarking systems, such as Delicious, tagging patterns tended to stabilise over time. New terms would tend not to appear and current tags would be repeated in roughly equal proportions at between 50 and 100 bookmarkings of an individual URL.

Bookmarking it is not a standalone practice, or indeed, product. This integration has to happen at the point of creating bookmarks and also at the point of use. If bookmarking is embedded in work practices (such as emailing, blogging and document creation) then it stands more chance of being used. If it is connected to search applications, personal profiles and even tools, such as Google Maps, then its collective benefits can be shared. This is best achieved when users no longer perceive bookmarking is happening. As the title of this article indicates, bookmarking will be successful when it disappears.

NEXT STEPS

If you want to introduce collaborative bookmarking to your organisation, the following advice may be helpful:

- Pick a tool that is easy to use and integrates with your other information technologies – eg email, search, document management, intranet, wikis and blogs. This may be an open source tool that you can customise or a commercial tool provided by a vendor.
- Start with something simple.
- Trial it with a group that has an obvious need for such a tool – ie they need to share information with each other, they use web-based tools frequently, they use a wide range of internal and external sources and they may not work in the same office. Examples of such groups are librarians, market and competing intelligence units or new product development teams.
- Make it easy for users to load their bookmarks from their browser or their Delicious account. Aim to get a critical mass of bookmarks and tags.
- Promote the tool and its benefits well. Recognise champions and power users.
- Be prepared for unexpected uses and benefits. You may have an information retrieval tool that is also an expertise location tool or even a map of the organisation’s culture.
- Develop a release schedule with upgrades to functionality and integration with other tools. Not all your experiments will be successful so work on the basis of “safe-fail”.
- Do not be disappointed if you reach a stage where the tool is used but no one sees it as an independent “thing” separate from your information environment. Disappearance (also known as ubiquity) can be a good thing.

¹ Golder S and Huberman BA, “Usage Patterns of Collaborative Tagging Systems” (2006) 32(2) *Journal of Information Science* 198.