



JISC cetis

centre for educational technology
& interoperability standards

JISC CETIS 2010 Informal Horizon Scan

Text © 2010 University of Bolton



This work is licensed under the Creative Commons Attribution 2.0 UK: England & Wales License.

JISC CETIS 2010 Informal Horizon Scan

Produced by various members of JISC CETIS in March 2010.

This report outlines some technology trends and issues of interest and relevance to CETIS. It should be seen as a set of un-processed perceptions rather than the product of a formal process; a great deal of ground is not scanned in this paper and it should be understood that no formal prioritisation process was undertaken. We hope it will stimulate discussion and recognise that this kind of material is by its nature contestable.

1. CURRENT TRENDS AND THE HERE-AND-NOW

1.1. Linked Data and the Low-fibre Semantic Web

Talk of the Semantic Web has for many years led most people outside a group of devotees to be either dismissive – “it will never take off in the real world” – or bemused. However, **uptake** of a digestible incarnation of the Semantic Web labelled “Linked Data” offers strong evidence that change is upon us both for Semantic Web adoption in the “real world” and as a herald of changed thinking about interoperability more broadly. We have seen what is claimed to be the first native Linked Data based educational application, Talis Aspire¹, large scale expression of public sector data² as Linked Data and a gradual evolution of usable developer toolkits.

From a theoretical point of view, the proposition is that the Linked Data approach standardises: the way data is written (i.e. syntax- the rdf bit); the way it is transported (its protocol- http); and the way you can identify concepts and other resources (identifiers- URIs). Some other combination of technologies could have achieved a similar result, but none appear to have the traction that the Linked Data technologies have. It is significant that the more elaborate components of the Semantic Web stack are consigned to the “only use these if you have to” bucket and that many of them stay there.

The result is that the time it takes from the intention to interoperate to working implementation is drastically reduced. In many cases, data integration is a matter of a tweak to a configuration file or a query. Nor does this approach require a wholly new set of systems. There are now fairly mature open source toolkits that can expose existing databases as linked data rather quickly. The combination means that, for the purpose of data manipulation, it is now much easier to look at what data is already available within and across institutions, and integrate it in real time.

From a policy point of view, the continued roll out of data.gov.uk (the government's Linked Open Data initiative) is providing more visibility, tools, and - most importantly - data. It will also lead to demands for JISC to make its data available as Linked

¹ <http://www.talis.com/aspire/>

² For central government, see <http://data.gov.uk/>

Data. There may be a conflict brewing here for education, which has been more concerned about the agendas of control and privacy, especially around data relating to students and university/college operations.

So far, we have largely seen interest in Linked Data be directed towards public sector statistics, research data and “Information Environment” concerns. We expect to see a flourish of creativity in educational applications but also note that the business world is making increasing use of Semantic Web technologies for integration, Enterprise Architecture, business intelligence...

The Educause Horizon Report 2009 mentioned “Semantic-Aware Applications” in the 4+ year timeframe, where the emphasis seems to be on latent semantic analysis rather than the semantic web technologies behind the Linked Data phenomenon.

1.2. Learning Platform Architectures

Distributed Learning environments featured in the March 2009 Horizon Scan. Since then we have seen a greater number of examples of a range of Learning Platform Architectures. For example, Google Wave and Opera Unite have attracted a lot of interest and provided good reference points for discussion. There are now enough examples (i.e. specific technologies) of alternative models that we can probably engage a wider range of people than would engage with purely abstract architectural considerations.

As outlined in the newly released CETIS briefing on Distributed Learning Environments³, five broad architectures are emerging for non-monolithic learning environments:

- One system in the cloud, that can push functionality into many environments, including an existing VLE
- Plug-ins to existing VLEs
- Web-based mash-ups
- Email and Google Wave style federations of servers with a choice of clients
- Platforms that both provide and consume widgets

Discussion

Of these, two and three have been around longest, but two appears to be falling out of favour, presumably because the cost of integration and maintenance don't outweigh the benefits for most people.

Three continues to improve gradually, but is still somewhat hampered by a lack of reliability and ID management infrastructure across platforms and tools. Since that issue is web-wide, it seems likely to be solved by a combination of OpenID and OAuth, or successors to these technologies. Once that happens, the current trend for educators to set up their own learning environments on the web could well increase considerably. For institutions, this means that looking at the ID management developments on the web could well be the main way in which some track can be kept at who is doing what, where. For the same reason, CETIS is starting to look at APIs for tracking learning interaction across various tools.

³ http://wiki.cetis.ac.uk/images/6/6c/Distributed_Learning.pdf

Models four and five still look somewhat more tentative. In the case of four, that's probably because the one modern product that could have pushed the federated server model - Google Wave - has faltered somewhat. It's a bit early to write off a technology that has so many interesting features, but discussions suggest that it may have to do with the lack of chronology in the user interface, which is the main organising principle in all other messaging platforms to date. Another factor could be the fact that Wave tried to support so many different types of task at the same time. As a result, two things could happen to Wave: either the user interface of the client is radically redesigned, or else the various technologies could wander to existing Google apps such as Google docs and Mail.

Opera Unite proposed an even more radical federation of servers, in that everyone with a browser client would also have a proxy server on their widget server. While it made that architecture much easier to set up and use, it didn't solve the perennial problem of how to deal with servers that are off the net as often as they're on it.

Lastly, the widget server that's also a client is still most strongly represented in Sakai 3, which is coming along nicely, and could provide a very smooth way into a DLE for adopters. Commercial competitors could well adopt the same architecture for the same reasons.

The first model, however, currently looks like the dominant solution with at least two commercial offerings specifically for the educational market (plugjam and icodeon CC platform), and two open source ones (shindig and wookie) next to countless general purpose widget platforms on the wider web that are more restricted in functionality (iGoogle, pageflakes, netvibe). The main reason for the traction that this model has is probably down to the simple facts that they allow adopters to keep their current VLEs and other webplatforms, and augmenting it with widgets is fairly cost and risk free.

A consequence of the wider adoption of model 1 might be that it could accelerate the trend to use generic content management systems such as Sharepoint as VLEs (see below), since the requisite educational functionality can be brought in in other ways.

The Educause Horizon Report 2009 mentioned "The Personal Web" in the 2-3 year timeframe. While this notes the rise of various component technologies, our view is that selective application of them in an institutional setting is likely over the same period as the "distributed Learning Environment".

1.3. Competence and learning outcomes

This topic was included in the March 2009 Scan but has proved to be somewhat more difficult to get a good grip on that anticipated. Perhaps the reality is that it is even more complex and significant that we thought. What remains at least as true now as last year:

- More and more people understand the potential value of being able to map pathways based on achieved skills and competence, both within education and from education to work.
- There is ever-increasing focus on learning outcomes in education, and therefore increasing demand for tools that can help manage these.

But where is a good place to intervene technologically? What are the "killer app" services that could be deployed? These are much less clear, particularly in view of the limited success of initiatives like GMSA Advance.

But perhaps it is precisely for these reasons that CETIS should keep a very keen eye on this area. We want to be among the first to identify and explore the realistic helpful application of technology to these challenges. We are continuing with the

valuable preparatory work, involving XCRI, Leap2A, HEAR, Europass Diploma and Certificate Supplements. We need to make sure we are attuned to signs of any ideas really taking hold.

1.4. Mobile Platforms Converging

The success of the iPhone App Store has resonated with the wider mobile industry, and there has been a rush to standardize many of the previously proprietary aspects of mobile devices, including the application platform and access to device capabilities such as address books, cameras and so on. There have been high profile partnerships based on a shared app store platform using the W3C Widgets specification, such as the Joint Innovation Lab and the Wholesale App Community. The focus has shifted away from the traditional low-level development models for mobile applications (e.g., Java, SDKs etc) towards HTML5, JavaScript, CSS and the web as platform.

This trend should become more apparent as more handsets enter the market, and as W3C specifications start to reach recommended status later in 2010. We should expect to see some of the innovations that emerged as a result of the introduction of the iPhone and App Store spreading to many other platforms and devices, including computers and TVs. The iPhone has 100,000s of applications despite a limited user base and high barrier to entry for developers - with a web-standards based mobile platform and a market of over 1bn handhelds, we can expect even greater innovation.

The Educause Horizon Report 2009 devoted one of its six principal sections to "Mobiles" in the <1year timeframe. The broad trend is undisputable but it is the undercurrents, consequences of shifts in business models and foci of competition that harbour both greater interest and greater uncertainty.

1.5. Augmented Reality Going Mobile and Mainstream

Augmented Reality appears to be moving into the mainstream⁴ thanks to smartphones with a combination of location awareness, direction awareness, and cameras. This is relatively lo-fi AR - no one is wearing funny goggles, just waving their phone around - but could catch on very quickly as an accepted form of user experience rather than just a gimmick. There might be a parallel to draw with the success of the Wii in spite of its relatively lo-fi graphics.

The Educause Horizon Report 2009 mentioned alternative reality in the 2-3 year timescale under the heading "Geo Everything". Geo-location capabilities seem to be dominant in current AR that is reaching the consumer but we anticipate applications in education and training that use local spatial markers⁵ coming out of the R&D phase in due course.

1.6. Sharepoint

Its old, its boring, its clunky, it doesn't work properly on anything but IE, but it is getting more and more deeply embedded in institutional IT. And as more institutions move to open-source LMSs like Moodle, more traditional portal functions head back

⁴ There are plenty of demo examples on the web, e.g. <http://blog.pachube.com/2009/06/pachube-augmented-reality-demo-with.html>

⁵ E.g. <http://www.gizmodo.com.au/2009/09/bmws-augmented-reality-glasses-turns-man-into-master-mechanic/>

to Sharepoint. In fact, the quiet trend in FE and training is to make Sharepoint the VLE, either by adding on a dedicated "LMS web part" or just cobbling together the bits needed. Its a bit old-skool, but its definitely a growing trend⁶.

Sharepoint could be/become the elephant in the room re standards/interoperability: "we have share Sharepoint, let's just use it (for everything)". There are some interesting examples of webservices being built around it beginning to emerge and we probably do need to think about how we can help share community knowledge around extending use of Sharepoint.

1.7. e-Textbooks

eTextbooks are not a new idea by any means but are one that has reappeared recently in the context of Open Educational Resources. Interest in eBooks in general appears to be stimulated by, or feeding back into, a growth in eBook readers and lately more broad support of ePub (Open eBook) including in the iPad.

Most OER activities are currently focused on the release of content however the next step is how to promote and facilitate the use of such content. With ever increasing financial constraints, institutions are perhaps also beginning to consider how they can use existing resource more effectively. In this context e-textbooks created from open educational resources may be seen as worth of further exploration. To date there is very little (institutional) practice in this area but the concept of OER e-textbooks is being raised in discussions with increasing frequency.

1.8. Image Analysis and Semantic Query

Recognition of faces and objects in images is fast becoming mainstream as content-based image retrieval matures and starts to be combined with ontology-based query interpretation. For example, Picassa⁷ can group similar faces. Imense⁸ can search, for example, for images of "two puppies in the snow" with good precision and their query-by-example, where images of the same or similar things are identified, is impressively accurate.

There are an immediate commercial applications to stock photography brokerage, to social networking to find lost friends and for "reputation management", to military and police intelligence or simply making the vast supply of user-generated images discoverable but what about education? One opportunity is for audio-description of images for visually-impaired or blind people, but the performance required for this may be some years away.

1.9. Rising interest in "Business Intelligence"

The label "business intelligence" (BI) has, of course, been around for a long time and often been misused for performance reports that gloss-over the detail necessary for decision-making. We mentioned BI in the March 2009 "Horizon Scan" and there seems to be growing interest in the Higher Education sector, probably driven by continued financial difficulties and uncertainty over the macro-forces influencing the nature of the business of HE. The fallout from the Spellings Commission in the US may also have been an influence in the rise of interest in BI through the marketing activities of international IT

⁶ Eduserv report <http://www.eduserv.org.uk/research/studies/sharepoint2009>

⁷ <http://picasa.google.com/support/bin/answer.py?answer=156272>

⁸ <http://www.imense.com/>

suppliers. Anecdote suggests that untested hypothesis, convention and dogma are used as the basis for decisions more often than they should be; there seems to be a need for expanding the application of BI.

If it were simply a case of HEIs adopting COTS and general practice in the use of BI, the above observation would not be worthy of a CETIS Horizon Scan. Four areas of enquiry appear to be worth investigation:

- semantic web technologies are beginning to be applied and the idea of "semantic middleware" is emerging to accommodate the sprawl of sometimes-subjective data in an organisation
- the character of the business of HEIs, their internal organisation and values differ from the mainstream commercial world and we should be cautious of mis-applying exogenous methods
- different realisations of BI could be applied through the whole vertical range of a HEI from senior management down to students. E.g. the techniques of BI might be usefully applied to students making decisions about their learning etc
- common metrics and information models for input-data rather than ad-hoc solutions could offer a sector-wide benefit .

2. ALMOST HERE

2.1. HTML 5 and next-gen browsers

HTML 5 introduces a number of innovations that are being rolled out in the latest browsers (Firefox 3.5, Safari 4, Opera 10... and lately the in-development IE9). These innovations include connecting microdata (similar to RDFa) in pages to drag and drop user interfaces, enabling new types of interoperability - for example, it is possible to drag and drop vCards for people between websites.

This interactive form of interoperability may become an important means of connecting applications at the user level rather than through back-end provisioning. Putting the user in the data exchange should help to avoid the "one size fits no-one" syndrome where users have to put up with what the software does to them. It could also significantly reduce the complexity in the design and implementation of interoperability specifications by reducing the number of semantically-well-defined elements that are required along with processing rules to use them. At the risk of stating platitudes: "users" have un-encoded knowledge that they should be able to exercise with the IT there to amplify their actions.

The thought of "putting the user in the data exchange" leads to adverse reactions from some people; we need to prove good examples to demonstrate the principle beyond pointing to examples in their MS/Mac/Linux desktop. How would a MIS/VLE "integration" look if, rather than exchanging the entire structure (modified by curious and arcane mappings) with IMS Enterprise, tutors could associate a VLE site with a module or course by dragging an icon for it from the MIS and then assign groups by the same method, leaving the back-end system to maintain updates only.

2.2. Open Data and the need for a new information literacy

While teachers, librarians and policy-makers are still grappling with information literacy in the era of the web of pages, we may be about to see a new twist develop. If the current surge in openly publishing both research and public sector statistical data continues, we will have a wealth of raw material for in-silico experimentation and an explosion of software to manipulate it. A rising need for inter-disciplinary approaches to problems ranging from social care to sustainable industry can be supported by this new resource but there is a problem: differing conventions and models in the different communities of practice. This suggests we need to develop new forms of information literacy in Higher Education that go substantially deeper than information discovery and provenance-assessment. It suggests that we may need to make develop in students the capability to read meaning from the conventions and implicit models of specialist disciplines other than their own. This will be challenging.

3. MORE SPECULATIVE DEVELOPMENTS AND INTERESTING OUTLIERS

3.1. Non-relational Databases

There's a quiet revolution going on in storage with the need for many systems to store really, really large numbers of simple name-object pairs in ways an RDBMS can't handle effectively. Apache CouchDB, Hadoop, Voldemort are all projects under the "NoSQL" banner⁹ that are trading query complexity for outright performance - almost the polar opposite of triplestores and SPARQL.

The result could be that the very wide spectrum of use cases for which relational databases have traditionally been used will get narrower. On the one hand, the NoSQL approach could start to become the solution of choice for cases where speed is everything, concurrency high, and data sets are large, homogeneous and run on a single (cloud based) system. On the other hand, Semantic Web approaches could encroach on cases where queries are analytical.

3.2. Real-time services

A recent trend is towards reducing the latency of asynchronous services to real-time or near-real-time updates. Examples include Twitter Search, Google Wave, and RSS Cloud. These reuse existing technologies such as Comet and XMPP (extensible messaging and presence protocol), and in some ways recall the early ideas of "push" web services.

As new services are designed, there will be greater user expectations of real-time or near-real-time updates. Any services in the HE/FE sectors should be taking note of this trend, and new applications being developed for the sector should be considering using real-time features.

3.3. Declarative Workflow

Conventional workflow focuses on prescribing permitted activity and model-driven software engineering from Business Process Modeling Notation (BPMN) leads us down a similar path. And yet, many of the activities in a HEI are subject to significant elements of judgement.

"Declarative workflow" or "constraint-based workflow" is an emerging alternative approach, although neither term is defined in wikipedia and "declarative workflow" has been used for approaches where an XML file is fed to a conventional-paradigm workflow engine (which is not what this section is about). The idea of declarative workflow is that only those sequences that cannot or should-not occur are defined. "Cannot" might be for data integrity, logical or "hard policy" reasons. "Should-not" captures good practice and advisory pathways. There is some particularly interesting work from Eindhoven University of Technology¹⁰ employing a formal constraint language and with prototype software. Work may be enacted in any way consistent with the constraints and data mining employed to discover actual pathways and social interactions.

Aside from application to administrative tasks, it is interesting to speculate about how the declarative workflow approach (and back-end software) might be applicable to teaching and learning in HE. Could it offer flexible support for structuring individual

⁹ <http://blog.oskarsson.nu/2009/06/nosql-debrief.html>

¹⁰ van der Aalst, Pesic and Schonenberg, <http://www.win.tue.nl/declare/research.html>

and group learning activities while allowing for individuality and autonomy? Does a focus on constraints, especially good-practice, simplify the task of defining the structure compared to orchestrating the whole process?

3.4. Dynamical Modelling of the Enterprise

The commercial world is beginning to develop general methods of modelling the enterprise that take account of the reality that the business and customers comprise a complex adaptive system¹¹ and that transformation to a desired "to be" state must take account of the response of the system to imposed changes. Of particular interest is the extension of the toolkit to include agent based models¹² (ABM) which help with understanding hard-to-analyse "emergent" behaviour (i.e. patterns in the system that arise from connections between components) given an understanding of the behaviour of individual agents.

Although ABM studies are probably not affordable at an individual institution, they may be viable at sector scale. Subject to their being sufficient data to validate models, it might be possible to better understand and predict the behaviour of staff and students (throughout the whole relationship cycle) that might arise from different modes of delivery and new applications of technology.

4. COPYRIGHT AND LICENCE

This work is © CETIS 2010 and licensed for use under Creative Commons Attribution terms: [Attribution 2.0 UK: England & Wales](#) and [Attribution 2.5 UK: Scotland](#).

¹¹ http://www.pwc.com/en_US/us/technology-forecast/winter2010/interview-mark-paich.jhtml

¹² http://en.wikipedia.org/wiki/Agent-based_model

About this Paper

Title: JISC CETIS 2010 Informal Horizon Scan

Authors: Various members of JISC CETIS

Date: April 2010

URI:

Text Copyright © 2010 University of Bolton; cover image courtesy of JISC



This work is licensed under the Creative Commons Attribution 2.0 UK: England & Wales Licence. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/2.0/uk/> or send a letter to Creative Commons, 171 Second Street, Suite 300, San Francisco, California 94105, USA.

For more information on the JISC CETIS publication policy see http://wiki.cetis.ac.uk/JISC_CETIS_Publication_Policy

About JISC CETIS

JISC CETIS is an Innovation Support Centre for UK Higher and Post-16 Education sectors funded by the Joint Information Systems Committee (JISC), and managed by the University of Bolton. The Centre provides strategic advice to JISC, supports its development programmes, represents the sector on international standardisation bodies and works with the educational community to facilitate the use of standards-based e-learning.

For more information visit our website: <http://jisc.cetis.ac.uk/>