Digimap: 10 Years Old and yet Brand New. The migration of an established service to new GIS infrastructure

Tim Riley, Service Development Project Manager, EDINA, University of Edinburgh

Abstract

For 10 years, the EDINA Digimap service has provided online mapping and geodata download tools to UK Tertiary Education. Currently, Digimap serves 50,000 active users in over 150 Educational Institutions. In order to provide appropriate performance, scalability and interoperability to meet increasing demand, it was necessary to replace the core GIS around which Digimap is built. So started a process of procurement, followed by a major reengineering project to implement a combined Cadcorp SIS® and Open Source (OSGeo) based infrastructure, while in parallel, maintaining an uninterrupted 24/7 service. EDINA and Cadcorp established a co-operative working relationship, which has benefited both parties and most importantly, the users of the Digimap Service.

Introduction

January 2010 saw the tenth birthday of EDINA[1] Digimap®, which since 10th Jan 2000 has been delivering on-line mapping and geographic data download tools to the academic community, 24/7, 365 days a year. During this time there has been continual growth in its use: it currently serves 50,000 active users in over 150 educational institutions. There has also been growth in the amount of data provided: Digimap began as the single point for the academic community to gain access to National[2] coverage of Ordnance Survey (OS) data but now also provides Historical OS (Landmark), Marine (SeaZone), and Geological (BGS) data.

The GIS on which the service was built did a great job but in 2004, the vendor ceased its development. Consequently, the performance and scalability of the service infrastructure was frozen, while usage continued to increase, as did the need to provide new tools that met the expectations of users: this at a time when the geo-web was rapidly establishing itself.

This saw the start of a long journey: from the Business Case to secure funding, through procurement, installation and rigorous testing of the chosen solution, reengineering of the Digimap platform and finally, delivery of new web applications, all without interruption to our 24/7 service. This paper tells the story of that journey and aims to show how each stage has lead to a robust new service architecture that combines commercial software with open source software and standards.

The Problem

In 1999 EDINA procured the Laser-Scan (now 1Spatial)[3] GIS software (Gothic, Lamps2) to provide the mapping services for Digimap. Back then, just as now, the cartographic quality of onscreen and printed maps was paramount and the Laser-Scan software met EDINA's stringent

[1] EDINA is a JISC funded National Data Centre based at Edinburgh University, whose mission and purpose is to enhance the productivity of research, learning and teaching across universities, research institutes and colleges in the UK. http://edina.ac.uk (EDINA, 2010)

[2] Great Britain

[3] This paper "Laser-Scan" rather than "1Spatial" because of the historical context.

requirements. EDINA developed the Laser-Scan software and a middleware application layer to provide a range of web-map browsing and large format hardcopy plotting services.

By 2004, Laser-Scan had discontinued development of their Gothic mapping software. So, while Digimap usage continued to grow (Figure 1), the underlying service architecture remained static. Over time, this caused increasing operational problems, with key staff spending 60% of their time ensuring the system stayed up and running. The system could handle only a small number of concurrent map requests, leading to requests being queued. This was not a problem initially but as usage increased, the queues got bigger, and Gothic increasingly felt the pressure. By 2006, parts of the service would need restarting on a weekly basis and all restarts were manual. Despite these issues and because of the team's technical firefighting, the Service remained up and running, with very few significant breaks in service. However, there were two other key issues:

- with increasing usage, performance was beginning to suffer and
- incompatibilities between the existing software and new web-based spatial solutions, meant that Digimap, one of the first web-mapping services, was now looking a little dated.

Clearly, the only way Digimap performance, stability and functionality could be improved to provide a future-proof quality service, was to build a new infrastructure, based on OGC[4] standards, OSGeo[5] software, with an innovative GIS at its core.

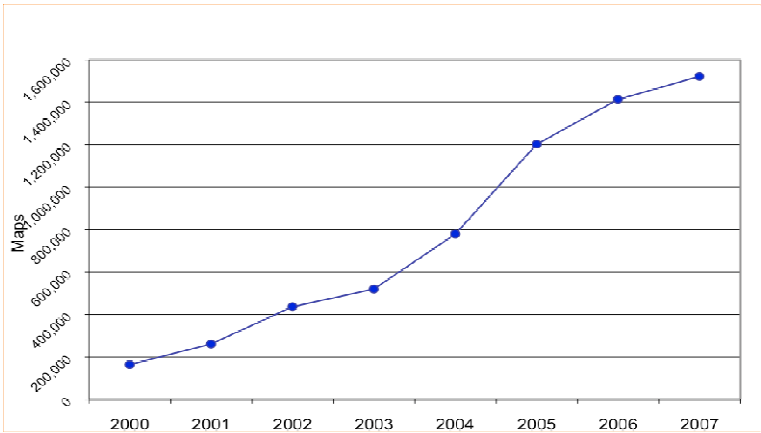

Figure 1: Number of Digimap "Classic" on-screen maps drawn 2000-2007.

The Procurement Project

The Business Case for funding

Digimap is funded by Central Government, via the Joint Information Systems Committee (JISC). The building of a new Digimap architecture would only be possible if JISC could provide the funding. JISC provides funding for many services and projects and all are important to Tertiary Education. Therefore it was vital that a clear Business Case for Digimap re-engineering was presented to JISC. So, in the summer 2006, *Informed Solutions* began a requirements investigation, assessing how the community was using Digimap and what was required to ensure service delivery continued at acceptable levels. Such an independent, objective assessment of requirements was essential, to validate EDINA's concerns and add weight to the Business Case. The findings of this assessment would also inform the requirements set out in the *Invitation To Tender* (ITT).

[4] Open Geospatial Consortium (http://www.opengeospatial.org)
[5] Open Source Geospatial Foundation (http://www.osgeo.org)

The Procurement Process

With agreement from JISC that the procurement of new GIS software could begin, in July 2007 the replacement GIS ITT was advertised in the *Official Journal of the European Community* (OJEC). This was a year after the Requirements study began and in that time, usage continued to increase (Figure 1) and EDINA continued to keep the service running and maintain performance, despite the platform limitations.

Procurement was undertaken by EDINA, in collaboration with the University of Edinburgh's Procurement Office, using the Higher Education sector's *Whole Life Costing Tender Evaluation Model* [TEM][6] to ensure business objectives were met (Table 1). With European Union law relating to procurement changing on what seemed like a daily basis, the guidance provided by the Procurement Office throughout the project proved invaluable.

Improved service provision	To allow delivery of a high-class service that continues to meet user needs.
Improved service availability and reliability	To improve availability of the service, minimising downtime, service interruptions and rapid service recovery.
Improved data Management	Improved data management processes, saving staff time.
Improved operations	To minimize costs and time to implement and maintain services
Improved capacity management	To measure, monitor and tune the service and its components for optimum performance. To allow the service to scale to meet growth in demand and new types of usage. To support new or modified facilities and future workload.
Improving technical interoperability	Improved mechanisms, based on open standards, to allow the service to interoperate (as service provider and consumer) with other services.

Table 1: The Main Business Objectives of Digimap Reengineering Procurement

Suppliers were required to respond to the ITT by 10 September 2007. Eleven suppliers downloaded the ITT from the University e-Procurement site and four responded with bids.

The Selection Process

The selection process involved three stages: 1] an initial screening of bidders; 2] a short-listing of bidders and 3] a final selection of preferred supplier. To obtain a shortlist, bidders were evaluated using the following criteria (in no order):

- value for money in relation to functionality offered
- ability of solution to meet operational and business requirements, based on responses to the Technical Response
- company viability, using accepted financial measures e.g. current ratio and liquidity ratio
- engagement with Terms and Conditions.

Two bidders were short-listed following the evaluation, based on their written responses. Of course, a bidder could supply all the right answers in an ITT response but this is no proof that

[6] The Higher Education sector's Whole Life Costing Tender Evaluation Model [TEM] provides a systematic method for the preparation of tender documentation, in terms of identifying and structuring the various criteria. Then, once the tenders have been received, the model is used to assess the submissions in a fair, transparent and ordered manner. TEM uses a combination of the full (whole life) financial cost, fitness for purpose (technical specification), qualitative and environmental and social judgments to help an organisation reach the best value for money decision from the available tender submissions.

their claims can be substantiated in practice. Therefore, a key component of the final selection was a consultation process involving EDINA staff, foccussing on practical technical assessment.

First, members of the project team visited each shortlisted bidder. The bidders were asked to present their proposal and demonstrate their software with respect to a set of scenarios provided by EDINA (Table 2). Suitable test data and detailed scenario instructions were provided to the short-listed bidders three weeks before each visit.

Functionality	Task examples / demos requested
Data Preparation	Import of sample raster and vector data; import into user defined data model.
Geoprocessing and Georeferencing	Georeferencing to OSGB coordinate system; display raster data in alternative spatial reference systems, e.g. WGS84; Create polygon data from OS Strategi line data.
Cartographic Map Display	Display a map symbolised to match a sample provided; overlay vector over raster and show transparency in web browser; demonstrate symbol creation process;
Database Update	Multiple user editing; editing roll back; simple geometry generalisation
Search, Querying, and Spatial Analysis	Select by buffer operations; report generation; select individual features and show attributes; demo in GIS and web browser.
Development Environment	Use development API to create web app functionality; API capabilities; system performance reports.
Interoperability and Database Connectivity	Display vector data from a networked database; demonstrate OGS WMS, WFS functionality;
Data Export	Export as GeoTiff and view tags; export whole and subset data as Shapefile and KML

Table 2: Summary of items covered in Evaluation Scenarios

Second, four reference sites, two provided by each bidder, were visited by the project team. The visits included demonstrations of the software and question and answer sessions. The latter covered areas such as functionality, performance, quality of training and support, working relationship, ease of development and handling of problems by the bidding organisation.

This stage was completed in December 2007 and evaluation of shortlisted bidders began. Scores were assigned to categories derived from all the information gathered, e.g.

- company viability;
- quality of after-sales service and technical back-up;
- a measure of ability to form a good working relationship with the customer;
- a measure of the bidders (based on terms and conditions, long term viability, quality of service, pricing policies etc)
- The ITT response.

In turn, these were grouped and weighted as follows:

- Financial (whole life: 5 year span): 20%
- Technical: 50%
- Quality: 30%

The scores for each dimension were combined using their respective weighting to give an overall score and rank for each supplier. At this point the evaluation model showed a clear favourite and this supplier was requested to provide a best and final offer (BAFO) in early March 2008.

We Have A Winner

Cadcorp[7] SIS®[8] was selected as the most appropriate product and in April 2008, EDINA was able to make a recommendation to JISC and submit final costs for the reengineering of Digimap.

General terms and conditions were specified in the ITT and Cadcorp had suggested informal agreement to these in their response. However, in order to draw up a formal contract, lawyers for each party persevered to ensure all requirements were covered in unambiguous terms. Meanwhile, the other bidders were informed of the choice of supplier and a standstill period began to allow them to challenge the result. There was no challenge but the legal tweaking continued and it was not until July 2008, one year after the start of the procurement project, that the final contract between Cadcorp and EDINA was signed.

The time taken up by the procurement was justified in that the process had been robust, thorough[9] and produced a successful outcome. However, it was now 4 years since Laser-Scan (by this time, 1Spatial) had stopped developing the Gothic GIS that sat at the heart of Digimap. Still usage continued to rise, with no capacity for extending the existing architecture to improve performance and scalability. The development of a new service infrastructure was now imperative.

The Digimap Reengineering Project

With the hard work of procurement complete, all that was left to do was build a brand new Service architecture; oh, and to maintain the existing live Service while building its replacement. So that there were sufficient hands to ensure the current live service was maintained while the new system was built, two new software engineers and a new GIS Technician were employed. This gave more experienced staff time to focus on the reengineering. It was now that EDINA's working relationship with Cadcorp began to develop. The success of this relationship would be a key factor in building a new architecture, which provided the necessary performance, scalability and interoperability.

Things were challenging from the onset. The EDINA geo-team knew what they wanted to do but now had to understand how they would get Cadcorp SIS to do it. The product is very different from Gothic and Lamps2. To assist this process, Cadcorp provided good in-house training on all aspects of the product suite. In parallel with EDINA's learning curve, the Cadcorp SIS software had its own challenges: The Functional Acceptance Tests (FAT) and the Performance Acceptance Tests (PATs), which Cadcorp were contractually required to pass.

The Acceptance Tests

Defined during the procurement process, the Functional and Performance Acceptance Tests were vital to ensuring the procured software provided the flexibility, stability and performance required. It was important to run the tests early in the project so that there was time for Cadcorp to address any test failures.

The FATs were based on the Evaluation Scenarios EDINA produced for procurement supplier demonstrations (Table 2). Each defined task demonstrated a key item of functionality, which the software could pass or fail. Cadcorp SIS passed on all counts in September 2008. The PATs were more challenging. They used full national loads of selected map products[10] to test continuous usage (Endurance Tests) and heavy usage (Load and Stress Testing).

---

[7] Computer Aided Development Corporation Limited
[8] Cadcorp Spatial Information System® - see http://www.cadcorp.com/.
[9] Large public procurements such as this, must be robust in order that they are defendable against challenge.
[10] OS MasterMap Topographic; OS Strategi; OS 10k Raster; Landmark 1:10,560 County Series First Edition raster; BGS 1:50k geology product

Endurance testing involved prolonged usage of the software, (3 days continuously, repeated 3 times), to ensure stability after continued usage and that the software had no memory leaks. **Load and Stress testing** involved heavy use of the software over a short period of time to ensure that it did not break, crash or become totally unresponsive under extreme load and bursts of activity (simulating 300 users making concurrent map requests for 2 hours, repeated 10 times). For both tests, PNG images were requested through WMS (to imitate the delivery of on-screen maps) and hardcopy PDFs of a random area generated (as is done by users creating map output for printing).

The initial PATs, (Jan 2009) did their job and exposed stability and performance issues in the Cadcorp GeognoSIS ™[11] software. In his blog (Daly M A, 2009) Martin Daly, Technical Director of Cadcorp, discusses the technical issues Cadcorp needed to over come in order that GeognoSIS could pass the PATs, which it did in May 2009. Cadcorp were diligent in their response and ultimately, gained improvements in their products due to EDINA's thorough testing regime.

Cartographic Design

The cartographic quality of on-screen and printed maps from Digimap has always been of paramount importance. Gothic software was powerful in this respect and allowed for fine control over all features and symbolisation. So cartographically, Cadcorp SIS had a hard act to follow. In one aspect, it won outright, as it allowed for high quality PDF map output for printing, a big improvement on the EPS output from Gothic. However, there were some issues and Cadcorp were quick to improve their software to meet EDINA's requirements. For example:

- stroke effects were added to drawing tools, vital for clear cartographic design.
- Interleaved drawing of multi-level pens and symbols were improved to allow clear road casing at roundabouts and junctions (where multiple features converge but only one master outline is required).

Also, the GeognoSIS WMS[12] interface did not provide fine enough layer control to allow EDINA to manipulate feature rich products (e.g. OS Strategi) as required. This was not something Cadcorp could change in the time available, so EDINA built its own application ("Clive" in Figure 2), which wrapped GeognoSIS with a WMS interface. This provided the extra functionality required to produce high quality customised maps and prints. Cadcorp are currently making improvements to their WMS interface to support the extra options needed for fine-grained layer control.

System Architecture

While work on the PATs and data management progressed, the new architecture was designed and implemented. The system is spread across a number of Solaris UNIX servers (databases, OGC WS, main web server, failover servers), and four Microsoft Windows Server boxes for the GeognoSIS Web Servers (three for live service load balancing, one for development). Within that hardware infrastructure sits all the software and data that constitutes "Digimap". The service is built from a combination of proprietary software (e.g. Cadcorp SIS, Snowflake GoLoader, GoPublisher) and Open Source geospatial solutions (e.g. TileCache, MapServer, PostGIS, Open Layers, MapFish). Figure 2 shows how these software elements relate and the flow of data through them.

[11] GeognoSIS ™ is Cadcorp's web mapping product
[12] Open Geospatial Consortium (OGC) compliant Web Map Server

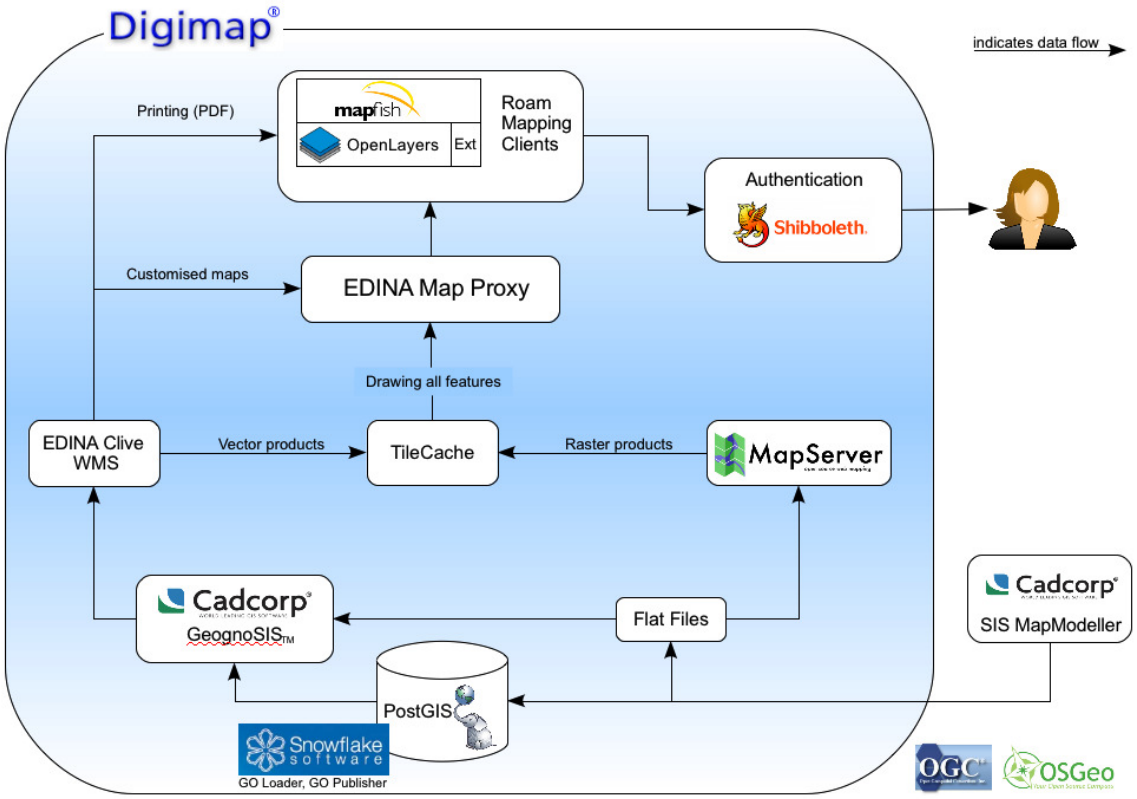

Figure 2: The new Digimap software architecture

Vector data are stored in PostGIS, with National OSMasterMap Topography coverage uploaded using Snowflake's GO Loader. Raster data are stored as flat files. In order to maximize web mapping performance, TileCache is populated with map tiles from GeognoSIS (vector data) and MapServer (Raster data), so that full GB coverage of all map products is cached. However, users can customise maps derived from vector products (e.g. turning off roads, buildings, etc) and it is not possible to cache all possible custom combinations. So, when a custom map is requested, it is delivered directly from GeognoSIS and the EDINA Clive WMS to the web-mapping client (Roam). GeognoSIS and Clive also deal directly with printing requests, generating high quality PDF output.

New Web Mapping Tools: *Roam*

The reengineering meant that new web map technologies could be included which the previous system could not support. EDINA could not only replace the core service architecture but also build new web mapping tools for users. Technologies such as OpenLayers were used to build a new web map application, *Roam*[13], for viewing and customising OS maps (Figure 3). The Roam framework brought Digimap's online map viewers bang up to date. It could now easily meet the expectations of new users familiar with the "slippy maps" of Google Maps and Bing, yet offered better quality GB cartography, printing and support materials than these other web-map tools.

[13] Roam: **R**apid **O**nline **A**ccess to **M**aps

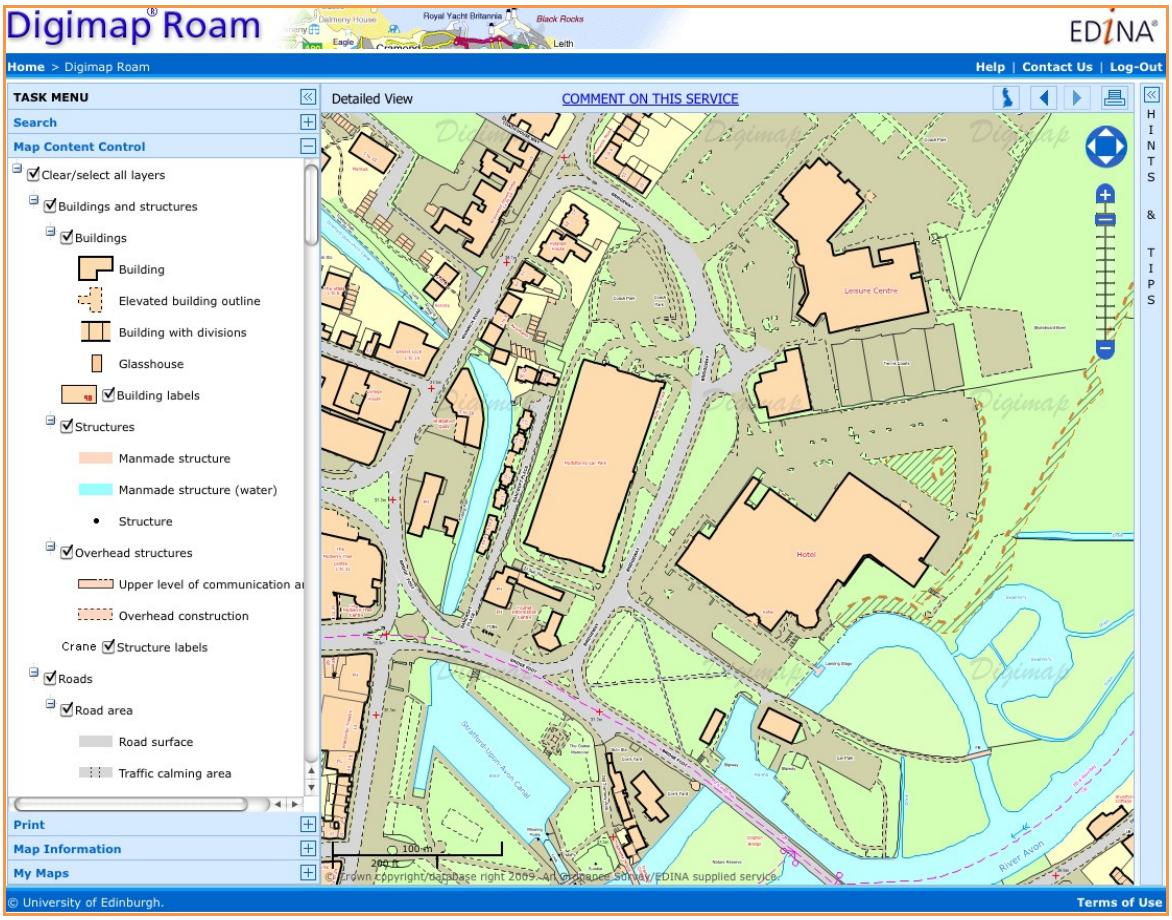

Figure 3: Digimap Roam, providing customisable online Ordnance Survey maps (©Crown copyright/database right 2010)

The New Service Launch

In September 2009 the new service architecture and the Roam mapping tool were launched. Other than seeing the new Roam application, the switch was invisible to users. Roam was well received by users, so, the same model has since been used to develop Ancient Roam (historical mapping) and Geology Roam, both released in the spring of 2010. Since the launch, up to June 2010, Digimap produced:

- 6,600,000 Screen Maps (1,867,843 were MasterMap)
- 170,000 PDF Maps
- 331,000 Downloaded Tiles of OS Data
- 650,000 km2 Downloaded of MasterMap Topography Data

Benefits

So, after all that work, what are the benefits? Well, here are just a few:

- EDINA can now focus on development instead of maintenance. From manually restarting processes on a weekly basis, restarts are now extremely rare. If they do occur, all is returned automatically.
- Stability. The service is very stable, with room for growth without the expectation of problems.
- Scalability. The system has capacity to deal with increasing usage and the introduction of new services and tools.

- Greatly improved performance. The old system could deal with only a few map requests at a time, queuing the rest. The new system can cope with multiple concurrent map requests and load balancing ensures the requests are passed to the most available server/processor.
- A basis for expansion and invention. As well as taking up too much of the team's time on maintenance, the old system was exactly that, old. Lack of performance and compatibility with new software meant that the service had no futurity. Now, Digimap has the scalability, performance and flexibility to allow for exciting new development to meet the needs of users.
- The users have a significantly more reliable service with better, faster mapping tools and hard copy print products.
- Cadcorp have benefited from improvements made to their software to meet EDINA's exacting requirements.
- EDINA and Cadcorp have built a mutually beneficial, productive relationship where knowledge on both sides has been shared and advanced.

Conclusion

Despite all the usual stumbling blocks, which conspire to bring complex projects down, from procurement to service delivery, the re-engineering of Digimap, has been a success on several levels:

- The procurement process was successfully managed by EDINA and the University of Edinburgh Procurement Office to ensure the selected solution met with stringent requirements. It is important to note that this selection needs to be based not only on the quality of the software but also the quality of the software supplier and their customer engagement.
- The procurement was a lengthy process but due to the associated legislative constraints, this is almost inevitable. It is vital therefore, that organisations embarking on such a process factor this into their planning.
- The re-engineering project successfully delivered a new service architecture and applications, which are stable, scalable and interoperable.
- The old service infrastructure was maintained throughout the new system build and migration from old to new was achieved without service interruption.
- The new applications improve a user's experience and productivity and have been well received by the Digimap community, which ultimately is what the whole process was for.
- The relationship between EDINA and Cadcorp has been positive, productive and mutually beneficial.

Digimap is now equipped for the present and the future, offering web based geospatial tools which will make a significant contribution to teaching and research as part of a developing academic Spatial Data Infrastructure (SDI).

References

Daly, M. *A Higher Education, Part 1*. Retrieved from Lost In Spatial: http://blog.lostinspatial.com/2009/12/09/a-higher-education-part-1/ (2009, December 9).
Daly, M. *A Higher Education, Part 2*. Retrieved from Lost In Spatial: http://blog.lostinspatial.com/2009/12/09/a-higher-education-part-2/ (2009, December 9)
Daly, M. *A Higher Education, Part 3*. Retrieved from Lost In Spatial: http://blog.lostinspatial.com/2009/12/09/a-higher-education-part-3/ (2009, December 10)
Daly, M. *A Higher Education, Part 4*. Retrieved from Lost In Spatial: http://blog.lostinspatial.com/2009/12/09/a-higher-education-part-4/ (2009, December 11)
EDINA *What EDINA Does: A Community Report. EDINA, University OF Edinburgh* (EDINA, University of Edinburgh 2010)