

ISC15D043

Title: *Update to website stability and people pages*
Author: Ben Petley/Paul Hooper (ISD)
Date: 6/6/2016
Circulation: ISSC – 14 June 2016
Agenda: ISC15A003
Version: Final
Status: Open

Issue

The following two papers were received by ET-P on 23 May 2016 and are supplied to ISSC as items for report.

Recommendation

Recipients are invited:

- To receive the reports

Attachments

- Website stability
- People pages

Project Title	Website Stability
Date	18 May 2016
Version	V1.0
Author	Ben Petley
Derived from template	Highlight Report last amended 11th February 2013
Overall RAG Status	Green

Summary of Progress

Background

- The University main website and portal are both delivered by a system called LifeRay.
- Since April 2015 there have been a number of outages of LifeRay that have meant the website and portal have been unavailable.
- The frequency of these outages has increased particularly during times of high load including over the start of year.
- Given the IT infrastructure and LifeRay implementation when hosted at UEA, outages required server restarts and resulted in unavailability of the website and portal for users for periods of 10-30 minutes at a time. Outages varied between several a day and once every few weeks.

Actions So Far

- Investigation and troubleshooting by CIS staff was unable to resolve the issues and as a result we engaged a consultant from Liferay to undertake an assessment and provide recommendations for how to resolve the problems.
- The consultant identified a range of issues which can be broadly categorised as follows:
 - bespoke developments that had been put in place to support specific UEA requirements
 - changes to product features that required recoding
 - insufficient awareness within the Web Team of recent Liferay developments
 - a need for better control of code
 - a need for better control of content creation
 - a need for better upgrade processes
 - a need for better system configuration
- Given the significant resource to carry out these steps and the efficiencies associated with moving to a hosted solution, the website was migrated to a LifeRay hosting company during the start of February 2016.
- All Liferay sites (main website, Portal, microsities) were migrated successfully to Firelay hosting at the start of February 2016.
- Initial migration failed due to too low an estimate of servers required to run our sites, with the issues we are experiencing.
- Through February and most of March end user experience was improved, although in no way perfect, as there were enough servers to mask issues and allow server restarts while other servers were still serving pages to users.
- DIT have carried out work to remove recommended content, configuration and theme, as specified by LifeRay consultants.

- Work has been carried out to amend database performance and caching of website content to improve performance of the system which should allow the reduction in the number of servers once all work has been deployed.
- Automated testing scripts have been developed to try and enable the stability issues to be replicated to aid in the identification of the cause(s) and subsequently necessary fixes. These have not been able to replicate the instability of the system, although will be useful in future deployment testing processes.
- All five members of the CIS Development team who carry out technical work on Liferay have completed seven days of advanced training.

Approach and Planned Actions

- At the guidance of the experts at the hosting and consultancy companies we are working with, the initial approach had been to try and replicate the issues seen in production.
- This has not been successful, so we are scaling back this work and concentrating on making best practice changes to improve the site.
- More formal project approach to be adopted, including project board, to ensure appropriate oversight.
- The project board will be chaired by Prof Jacqueline Collier and will further comprise Jonathan Colam-French, Angelina Bingley, Giles Whattam, Paul Hooper, Chris Hill and Ben Petley.
- To improve the efficiency of the running of the project, we have engaged an independent project manager who is a Liferay specialist, who will report to the project board, to bring the work of the hosting company, the consultancy company and UEA together to make it more organised and well planned.
- We have asked the project manager to put together contingency plans to mitigate risk to key University process, such as Open Day advertisement and Clearing, should the issues not be resolved by then.
- We have also asked the new project manager to investigate the approach for analysing the problem, in particular the requirement to introduce new code and functionality makes it almost impossible to troubleshoot problems as the system is constantly changing. There maybe a way to introduce new code into a separate system or server.
- A new review point agreed to be set to ensure that we have an adequate outcome in a reasonable timeframe. This review point is yet to be decided by the project board. The project manager has been asked to develop a plan B for implementation at this review point, should progress not be sufficient by then.
- A full project timeline, with planned changes and deadlines will shortly be put together by the new independent project manager and details of this can be circulated, once finalised.
- These changes will be made over 2 week periods and user experience metrics used to confirm improvements.
- A paper will be going to the next project board to approve the metrics which will indicate sufficient improvements in the user experience of the website.
- Although the migration improved the user experience by significantly reducing absolute downtime (i.e. users rarely experiencing “web page unavailable” messages), the page load times have increased due to changes required as part of the migration. The work to undo these changes is planned in, which should see an improvement in page load times and therefore user experience.
- Additionally, since the end of March the number of server restarts which the hosting company has had to undertake, has increased. This will have had a relatively minor impact on users and we are working with them and the software supplier to resolve these issues.
- Increased server count has meant increased ongoing costs (hosting and licences). Therefore the reduction of this server number has been included as another project deliverable.
- Improved governance procedures being worked up to support existing, ratified, policy.

- An intern position is being advertised to sit within the DIT (ARM) team, to carry out a number of manual content based changes which have been recommended to be made to improve performance.
- Deployment processes for new code have been amended to increase automation, automated testing and follow best practice for development. These are intended to stop any further instability being introduced in the future. These have been worked up with the consultants who are signing off the final process. Final details of this process are being agreed, prior to it becoming business as usual. UEA staff have already lost access to deploy new or revised code to the live sites (the hosting company do this now), which is enforcing best practice for deployments.
- The new deployment processes include the use of automated deployment and testing tools which are provided by our hosting company and have been configured by our consultants.
- The new deployment processes will include a regular release cycle and a paper presenting the options for this are to go to the next project board for decision and approval.
- We have engaged the consultants to carry out extra quality assurance/code review work on new developments (which we are currently trying to minimise) for the next few months.
- Improvements, combining and refactoring of code for various elements of the website will be completed by UEA and Worth.
- The deployment processes and systems will be completed and signed off by Worth.

Budget Position

Please see appendix A for the full budget.

All these monies have been found from existing ISD budgets and contingency.

Note that the budget will require some re-working as a result of the new approach.

Exceptions and Risks

- **Issue:** Firelay have reported that number of restarts on the website has increased since the end of March.
- **Mitigation:** Firelay and UEA have agreed a list of actions to investigate the possible causes which will allow us to resolve the issue. The two key areas being investigated are the new course catalogue (backend and frontend implementations) and implementation of Hotfixes for Liferay.
- **Risk:** Need a decision on how much time and money we are prepared to invest in the initial investigations on causes of the performance issue.
- **Mitigation:**
 - Metrics provide a measurement of impact of changes.
 - Detailed specifications detailing outputs required that we can measure.
 - Turn work into a full project which will detail in the plan what we expect to spend, timescales and what we expect to get from the work at the end including ongoing benefits defined.

- **Risk:** adding new content and coding to the live website makes it difficult to assess if the changes being made to improve stability are having the desired affect. New code may degrade liferay performance. By continually changing code it will prove difficult to put timescales to the discovery and anaylsis part of the project to resolve the stability issues.
 - **Mitigation:** ensure all new code goes through the quality assurance process and detailed change control is in place so we are aware of when changes are made. Appoint a project manager to clearly define the work and approach we should be taking to resolve the stability issues, have a plan in place for critical dates to ensure a stable environment, identify alternative options in case the existing work does not resolve the issue.

 - **Risk:** Lack of resource in DIT for testing following restructuring of codebase and to implement potential site restructuring
 - **Mitigation:** Recruit intern/temp staff to assist
-

CISB15D024

Appendix A - Budget

Website Stability/Hosting Costings						
Cost Type	Description	2015/6	2015/6 Paid	2015/6 unpaid	2016/7	2017/8
Licences	Previous Licence (2 x prod & 2 x non-prod @0.25; 13705)	£41,115	£41,115	£0	£41,115	£41,115
	New licences (1 x backup @ 0.50 & 1 x non-prod @ 0.25; 13705)	£6,167	£0	-£6,167	£12,335	£12,335
	Short term licences (1 x prod; 13705)	£0	£0	£0	£0	£0
Hosting	Original Server Hosting (1 x prod app & 1 x backup app & 3 X DB & 2 x acceptance & 1 x test)	£28,061	£28,061	£0	£42,091	£42,091
	Extra App Servers Hosting (2 x prod app)	£6,342	£2,114	£0	£6,342	£6,342
	Extra Web Servers Hosting (2 x prod web)	£2,861	£954	£0	£5,721	£5,721
	Firelay Worth Sandbox Hosting	£362	£362	£0	£0	£0
Migration	Firelay Migration	£18,727	£18,768	£41	£0	£0
	Firelay Server Setup	£8,752	£8,752	£0	£0	£0
Issue Resolution	Firelay Optimisation	£7,222	£2,417	-£4,804	£0	£0
	Revised Worth Discovery	£18,000	£0	-£18,000	0	£0
	Worth Consultancy 10 days & Blazemeter	£11,208	£0	-£11,208		
	Worth Backlog Implementation	£28,792	£0	-£28,792	£0	£0
	Firelay Tuning and Project Management	£19,455	£0	-£19,455	£0	£0
	Worth Code Review for Improved Deployment Processes (1 day a week for 2 months)	£8,160	£0	-£8,160	£0	£0
Training	Liferay Training	£32,868	£0	-£32,868	£0	£0
Ongoing	Worth Consultancy (2 days per month x 4)	£8,160	£0	-£8,160	£0	£0
Contingency Total		£184,364	N/A	N/A	N/A	N/A
Total		£246,252	£102,543	-£137,573	£107,604	£107,604

Title: **People Pages**
Author: Paul Hooper
Date: 16/03/16
Circulation: ISSC
Version: V1.0

Background

People information on research activity are currently presented on the main UEA website using the Liferay product.

The Digital Innovation Team (DiT) created a project to look at presenting key pieces of information relating to research activity at the start of 2012. CIS were asked to help assist in the technical aspects of the work whilst DiT consulted with stakeholders to construct the requirements of what the People Pages should do.

Two technical options were presented for discussion at the Web Steering Committee which included:

- Using the main UEA website product Liferay to present People information with a data interface into the research system known as PURE. PURE holds all of the key information required to be presented online.
- Purchase and setup the PURE Portal to present the information with a link from the main UEA website.

It was decided that the Liferay product should be used so as to provide the greatest level of flexibility in meeting UEA requirements. The People Pages were signed-off against the requirements defined during the second half of 2012 and deployed during the second half of 2013.

When People Pages were first made available DiT provided first line support to triage any problems that were reported, investigating if the problem was related to data within PURE/HR/ARCP system, to the website functionality or related to the requirements of what the system was asked to perform.

The main person within DiT providing first line support and had in depth knowledge of the design and requirements left the University in autumn 2014. This has resulted in reduced levels of support and questions being raised recently as to who actually owns and manages people pages and its data.

How People Pages are populated

Although PURE holds most of the relevant information needed for People Pages (email address, phone number and room number are fed from SPOT), PURE is fed key data from other sources.

- **HR/Payroll System to PURE:** Information sent to PURE includes the main person data and their post information such as school they are working in and role that they are performing.
- **ARCP to PURE:** Information is sent to PURE on research projects taking place, research applications and supervisor details.
- **SITS to PURE:** PGR supervision data is sent to PURE relating to which students specific academics supervise.

The data is presented on the main UEA website (Liferay) using a Liferay portlet which collects data from PURE. This portlet creates a cached object of data for each person containing all of the relevant information. When a user of the website clicks on a person with People pages the object data is presented from the website cache which speeds up the presentation of data on the web pages. Using this method prevents the need to interrogate PURE on a per click basis. Objects are cached by Liferay for a maximum of 12 hours and are then refreshed.

People Pages Availability

Users experience in trying to use People Pages has started to become inconsistent with some users finding the website unavailable and others finding individual people pages presenting an error message.

We believe that the problems being experienced are relating to two issues.

- **Instability of the main UEA website running on Liferay:** A separate paper has been presented covering the main issues being experienced and work taking place to resolve these.
- **Corrupted Cached Objects:** for some people pages we have recently found that the cache objects created by Liferay are becoming corrupted. This has the end result of producing an error when the data is presented. We believe that objects are becoming corrupted due to the size of data being presented.

Recent information from the hosted solution provider indicate that the servers are requiring a regular restart so as to clear the corrupted cached objects along with resolving other website instability issues. CIS are investigating with the external hosting company metrics that can be put in place to monitor the situation and provide quantitative information on the extent of the problem.

In order to resolve the cached objects issue, the web team are working on a short-term solution to improve the situation quickly and then assessing a longer-term resolution.

- **Short-Term:** add some logic to Liferay so that it tests for corrupted objects and if the object has been corrupted to clear the object and then recreate the object. Although this will allow data to be presented it may slow down the presentation of that data if the object has been found to be corrupt.
- **Long-Term:** perform an export of data from PURE each night to create a cache of the people pages in static web pages.

Additional Problems

DPA: Recently, ISD was asked if a PGR Student could have their details removed from the people pages. The original specification did not allow for individuals to not be included within the web pages. ISD have investigated the issue and put in place a workaround which at the moment means the student's details have been removed from PURE so as to prevent the data from being presented in the Website.

The PURE Portal which is the alternative to the People Pages in Liferay, does have the functionality to exclude individuals from the People Pages.

Multiple Identities: A recent requirement was to allow a person to appear in more than one location on the people pages. The People Pages presents information on researchers by school.

It is possible for someone to need to appear in more than one school or area if they have a current role specified by the HR system in that department. Due to how data is fed into PURE and therefore into the People Pages this does mean that the level of flexibility in where to present the data is limited. The limitation can be described using the Vice Chancellors record as an example of the problem:

- HR system provides the person data and which school and Tab (category) that they appear in
- HR system can allocate as many posts as required into the people pages
- VC has two entries, one as the VC and one in BIO under Honorary
- VC cannot appear in the academic section of BIO pages as he doesn't have an academic post in BIO anymore in the HR system.
- Workaround was to add the VC to the HR system as an Honorary post in Bio to prevent him being paid twice.
- Honorary post is a workaround because this is the only type of unpaid record available in the HR system
- Note: HR system is also the payroll system and for the integrity of the data in it HR must be careful as to which records they create.

CIS are investigating with HR if it is possible to add a new category that we can define in the people pages as being included in the Academic Tab but which is not a paid post. When using the PURE Portal we must ensure that this requirement is included in its development.

Reduced level of Support for People Pages:

Since the departure of the main person supporting People Pages within DiT issues have not been dealt with in succinct or efficient manner. This has provided users with a poor experience when reporting issues and attempting to obtain a resolution. Some users have had to talk to multiple groups in order to obtain a resolution.

RED have provided CIS with a sample selection of the type of issues that were frequently being reported to them. In almost all of the cases provided to CIS these related to data in PURE and could only be resolved by staff in RED. CIS have since provided RED with a spreadsheet detailing how they can investigate and resolve these types of issues.

Where issues did not relate to data, RED has been escalating these via CIS operations to the web team. Due to the current focus on improving the stability of the website, these issues have taken a very long time to investigate or resolve as the required resource within the web team was unavailable.

In order to improve the level of support being provided to users who are reporting issues, DiT, RED and CIS have agreed the following short term solution:

- RED have agreed to receive and triage all new People Portlet issues reported over the next four weeks.
- RED will record each issue reported and how it was resolved. This will inform and assist RED, DiT and CIS to define a more effective process for resolving issues going forward.
- Any issues that are not data related will be passed on to the web team who will prioritise these over the next four weeks.

This will help RED, DiT and CIS establish the number and types of problems being faced by academic and research staff when using the People Pages portlet. With this information it will be possible to look for resolutions to issues which are of high frequency and document methods used to investigate and resolve issues.

Summary

The PURE portal was on the CIS Roadmap to be implemented and was added to the CIS programme of work to implement as a project in 2015. Problems with ARCP and the data it was distributing to PURE meant that work on the PURE Portal needed to be stopped so as to resolve the ARCP problems.

The new Pure Portal provides an easy-to-use, searchable public facing website which features all of the university's research data including applications, projects, research outputs and researchers.

The Portal will enable us to achieve the following aims;

- Showcase all research activity and to stimulate both local and global collaborative partnerships and attract new research talent to the university.
- Help the university adhere to the new EPSRC statutory requirement that metadata on research data is available and visible externally.
- Improve the process for deposit of full text to increase overall compliance for REF eligibility and individual funder open access requirements. This will enable us to meet the HEFCE open access REF criteria.
- Simplify our research applications offering by reducing the number of interfaces required between REN systems.
- Easily identify like-minded researchers for strategic research collaboration
- Replacement of e-Prints our current digital repository with an application that meets RED and the Library's needs.
- Improve the appearance and accessibility of UEA's repository

We recommend that work starts to progress with the implementation of the PURE portal to help to resolve some of the issues we are seeing. This will mean that the uptime of the People Pages will be higher as they will not be hosted on the main UEA website and therefore will not suffer from the effects caused by instability within Liferay.

We recommend that this work is done as a project and included within this project should be the PURE upgrade and ref2020 work. This is so that the right sequence of events take place, interdependencies are known and documented and taken into consideration when the work is co-ordinated.

The key issue that will need to be resolved is the ownership of the People Pages and the PURE Portal. At this time no one group within UEA wants to take responsibility for the data or business process side of the PURE portal.