# SAKAI EVALUATION EXERCISE - Summary

**Rob Crouchley and Adrian Fish**
e-Science Centre of Excellence, University of Lancaster

**Rob Allan and Dharmesh Chohan**
CCLRC e-Science Centre, Daresbury Laboratory

e-Mail: r.crouchley@lancs.ac.uk or r.j.allan@dl.ac.uk

May 9, 2005

**Abstract**

This report summarises an evaluation of various aspects of the CHEF, Sakai and other related projects that may play a significant role in establishing a Virtual Research Environment (VRE) for the UK.

The work carried out built on e-Science activities at Lancaster and Daresbury, which include software development and hardware deployment for collaborators in the physical sciences, substantive e-Social Science research and social science training and awareness-raising. It built upon the existing collaborations that the project team have with each other and with other university groups in the UK. It is now using existing mechanisms to disseminate the outcomes of the work. Throughout the period of this evaluation we have received clear expressions of interest in working with us from the developers of Sakai (http://www.sakaiproject.org). This included discussions and presentations at the Sakai Developers Workshop, Denver, USA, 23-27/6/04 and subsequently at the JISC-CETIS workshop, Oxford 4-5/11/04.

The evaluation exercise was greatly facilitated by Lancaster deploying a CHEF implementation (http://redress.lancs.ac.uk) from day one and later the Sakai RC2 implementation to use as a collaboration tool for developers based at Lancaster and Daresbury.

This report summarises the outcomes and contains the conclusions of the evaluation which covered:

1. Comparing Sakai/ CHEF with Alternative Frameworks for VREs;
2. Establishing the ease of administration (EoA) of Sakai/ CHEF for a VRE;
3. Establishing the feasibility of making existing VRE (Grid) components available via Sakai/ CHEF;
4. Establishing the feasibility of extending the functionality of Sakai/ CHEF particularly to use Web Services for distributed development and deployment.

The full report is available on line at http://redress.lancs.ac.uk/sakai-uPortal and http://www.grids.ac.uk/Sakai/sakai_doc.pdf.

# 1 Introduction

Our primary aim is to evaluate the appropriateness of the Sakai/ CHEF frameworks for the development of Virtual Research Environments (VREs) within the UK education, digital information and research communities. This report should help JISC to formulate the long-term requirements of VREs. The relationship between, and capabilities of, the CHEF/ OGCE and Sakai frameworks is explained in the full report. Another JSR-168 compliant framework, GridSphere, is also described for completeness as it has a certain popularity for Grid projects in Europe.

A broad definition of a VRE was provided following some debate by the JISC JCSR VRE Working Group [1]. It encompasses some, but not all functionality of:

**e-Research:** access to research tools, personal information, project-related issues;

**e-Collaboration:** contact with and working with researchers in the same or related fields;

**Digital Information:** access to relevant resources to provide background and supporting evidence for research and training;

**e-Learning:** components relevant to awareness and training. We assume that more research-related learning will be self-motivated and guided by example rather than managed coursework. Assessment tools are therefore ignored;

**e-Management:** project and financial management in a research environment;

**e-Authoring:** authoring of all project and related research oriented material, often requiring input from partners and review by peers and chamions;

**e-Publishing:** publishing in appropriate format of research outcomes. Encompasses the peer-review process;

**e-Leisure:** BBC news, weather, current affairs, finance, local issues and events.

CHEF, the underlying technology which was the initial focus of this review, is an open source, freely available framework designed to provide useful collaboration tools, such as a chat facility, discussion boards, shared calendars and file sharing in the context of a course or tutorial 'worksite'. It currently provides analogues of the most popular features of commercial groupware offerings such as IBM's Lotus Notes. The Sakai project aims to bring CHEF and uPortal together and then add more VLE facilities. It is in fact designed to encourage open collaboration and the sharing and contribution of new tools by a large and growing community of users, and this appears to be appropriate to the needs of the UK research community. After some discussion it was felt appropriate to include GridSphere in the evaluation of open source tools as it is being used in some UK e-Science projects. Bodington was included for completeness because of its prominence as a UK Virtual Learning Environment.

Jetspeed and uPortal are open source, freely available portal servers that allow portal content to be imported, or linked to, by an administrator and then flexibly configured by users, thus empowering the users in the development of their own collaborative environment. They are hosting containers supporting portlet standards, but do not offer higher-level content management or tools

and therefore are not included in the final assessment. Other generic portal engines are listed in an Appendix.

Virtual Research Environments, VREs, by their very nature will continue to evolve. It is important to make sure that our existing resources, services and applications can be made accessible in emerging standards-compliant frameworks and can be supported in the long term. A part of the funding requested for this evaluation was for a UK subscription as early adopters to the Sakai educational Partners Programme, SEPP. There are obviously good financial reasons for moving to an open source platform for VRE delivery and collaboration, such as Sakai/ CHEF. The software is free and works on various platforms, being 100% Java code on the server side. There is thus no tie in to specific hardware, and the system will scale well financially with no licensing costs for extra server nodes etc. There are good logical and technical reasons also. Sakai will be completely open source and open architecture, so institutions will be able to customise existing, or add new, Java JSR-168 compliant portlet codes to connect to a legacy system running on their campus, e.g. to use existing project services such as databases, evaluations and timetables. Sakai will be configurable to use various different database management products, both commercial and open source, thus avoiding RDBMS lock in. At Daresbury an Oracle 9i RAC meta-data server and SRB are being used, but not all projects use Oracle so we tested PostgreSQL too. By exposing existing VRE services using the standard portlet API through a customisable portal framework, and re-using much existing code and sharing additional collaboration tools, we can achieve a large reduction in software development outlay and encourage closer community integration. Of course, this reduction has to be offset against the potentially increased need for software developers to throw away their bespoke solutions and adapt or interface to legacy systems so that they can co-exist with a VRE framework. If they do so however we can achieve a portable and maintainable solution.

This evaluation has addressed the following areas:

1. Comparing Sakai/ CHEF with Alternative Frameworks for a VRE. WP 1 is reported in Appendices A and B (the latter re-produced in the separate summary document);
2. Assessing the Ease of Administration (EoA) of Sakai/ CHEF for a VRE. WP 2 is reported in Appendix C;
3. Establishing the feasibility of making existing VRE (Grid) components available in Sakai/ CHEF. WP 3 is reported in Appendix D;
4. Establishing the issues involved in extending the functionality of Sakai/ CHEF particularly to use Web services for distributed development and deployment. WP 4 is reported in Appendix E;
5. Developing a *Roadmap for a UK Virtual Research Environment*. This is now in a separate report from the JCSR VRE Working Group [1].

The deliverables of the project summarised in the remainder of the full report report are:

1. Evaluation Report part 1: Technology Survey (WPs 1,2);
2. A Review of the Issues for Building Standards Compliant Portlets (WPs 3,4);

3. An Assessment of the Potential of Sakai/ CHEF as a Platform for Customised Portals, e.g. ReDReSS, NCeSS, e-HTPX, e-Minerals and NGS (WPs 3,4);
4. Evaluation Report part 2: Developer and User Feedback (WPs 2-4);
5. Software Template for Sakai/ CHEF Institutional Adapters (WP 2);
6. Software Template for Grid tool wrappers for use in Sakai/ CHEF (WPs 3,4).
7. *Roadmap for a UK Virtual Research Environment* [1]

Our work has been written up in this report and also included in a couple of papers to the 2004 e-Science All Hands workshop, see References [2, 3]. An interim report was submitted to Nicole Harris of JISC on 3/9/04.

Separate summaries, project deliverables and the full report are available from the *Sakai Evaluation* worksite of the ReDRESS Portal. It is accessible from `http://redress.lancs.ac.uk:8080/portal` by logging in with (username=guest, passwd=eResearch). The full report is also available from `http://www.grids.ac.uk/Sakai/sakai_doc.pdf`.

- *SAKAI EVALUATION EXERCISE* Rob Crouchley, Adrian Fish, Rob Allan and Dharmesh Chohan (this report)
- *SAKAI EVALUATION EXERCISE - Summary* Rob Crouchley, Adrian Fish, Rob Allan and Dharmesh Chohan
- *SEE: Architectural Summary* Rob Crouchley and Rob Allan
- *SEE: Technical Summary* Adrian Fish and Dharmesh Chohan
- *SEE: Interim Report* Rob Crouchley, Adrian Fish, Rob Allan and Dharmesh Chohan
- *Sakai Installation Guide* Adrian Fish
- *Sakai Architecture Evaluation* Adrian Fish

It is clear that, whilst we have attempted to be inclusive in our survey and review, such work is time-bound and never complete. We apologise if we have omitted to reference the full set of appropriate tools and technology and beg their authors to contact us with more information.

# 2 Technology Evaluation Report and Conclusions

## 2.1 Introduction

This work package compared Sakai/ CHEF with Alternative Frameworks for a VRE. By a "framework" we really imply a means to deliver the VRE services in an appropriate way to end users and developers. This comparison is restricted to the high-level open source and platform independent

frameworks, that is, Sakai, Gridsphere, CHEF/ OGCE, and Bodington. For the purposes of this report, we regard uPortal and Jetspeed as enabling technologies (portal containers) and do not review them here, indeed there are many others which fall into this class. This comparison is an ongoing task and lasted for the duration of the project as much of the technology and functionality is still changing.

## 2.2 Method

We have chosen to adopt a Multiattribute Utility Evaluation approach [4] to evaluating the alternative contenders for a VRE. This approach is nevertheless straightforward and simple to apply. The core of the procedure is to identify the most relevant values or criteria that are appropriate to the functioning of a VRE. Measurements are then made to determine the degree to which the criteria are attained. By doing so systematically, and by making numerical judgements wherever possible, we can compare the VRE contenders on a more objective basis than is usually the case. We have identified the following 10 broad criteria, some of which are further subdivided.

These criteria are as follows:

**Criterion A:** Can the framework be applied to all disciplines?

**Criterion B:** Is the VRE framework useful for users from both the e-science (a) and e-learning (b) domains?

**Criterion C:** Criterion C is divided up into two sub criteria:

a) It is generally accepted that software modules that are used by many projects end up being robust and well understood due to the amount of exposure they receive. It thus makes good sense to make use of publicly available libraries when building a software product as opposed to writing the same algorithms over and over again. This criterion is thus intended to reflect the degree of use of open source libraries by the VRE.

b) Conformance to ratified standards is another feature that is generally seen as being important. Standards conformance fosters ease of interoperability and ease of extension via plug-able components. JSR 168 and WSRP have been identified as the main two standards that, when adhered to, will allow Java components and Web service based components to be added to a VRE. What interface standardization also facilitates is reuse. There are many tools currently in circulation, written in Java, which could be re-factored to allow them to be plugged into a standards-based interface like JSR 168. There are also many tools written in other languages. As long as these tools can be re-factored to talk the WSRP protocol, then they could also be re-used in a WSRP compliant container.

**Criterion D:** Make UK services and resources available in familiar environments e.g. typically via a Web browser;

**Criterion E:** Any open source project worth its salt is supported by a decent online community. Discussion forums, chat rooms and mailing lists have all proven to be incredibly useful tools for spreading know how about a software product. The following two criteria are an attempt to measure this kind of support.

a) This measures the degree of support there is for developers who wish to start writing tools for, or even extending, the VRE framework.

b) This measures the degree of support there is for users and administrators who wish to install, configure and use the VRE framework.

**Criterion F:** Offer choice in presentation or delivery for (a) services and (b) tools;

**Criterion G:** How steep is the learning curve required to use the VRE framework?

**Criterion H:** This criterion measures the amount of functionality you get "out of the box" from a particular framework. In other words, if an institution installed the basic version of this VRE framework and just left it at that, would it be any use?

**Criterion I:** Presence and extent of a future funding stream, for (a) <12 , (b) >12 and <24 and (c) >24 months

Have all the criteria been listed? There are others we could include like the track record of the developer team, but this is taken into account when we allocate a score under sub criterion F(a). We felt that other, perhaps domain specific criteria, would be less important in an overall evaluation and are thus likely to have a lower impact and not affect the overall rankings. This has been partially tested with a sensitivity analysis (see below).

We then ranked the 9 criteria in order of importance and allocated a score out of 10, this was then standardized to sum to 1. If criteria had a lot of overlap with another criterion they would be given similar rankings. For simplicity, and to start with, we have given them all equal weights (0.1).

The same process was then applied to the sub criteria. This gives us the weights $w_{Xj}$ that can be used to aggregate the score of each component of criterion $X$ to produce a total score for criterion $X$. So for example, for criterion $B$ which is subdivided into 4 components, the total score for criterion $B$ is given by $B = \sum_j w_{Xj} B_j$. The total score over all the criteria is then given by $U = \sum_X r_X X$, where $r_X$ is the weight applied to criterion $X$. The weightings we used are shown in parentheses in the following figure.

The scores (out of 100) for each component were then obtained using our judgment.

## 2.3 Score Justification

### Criterion A

All the platforms tested are domain agnostic before specialisation; it is the tools that make them domain specific. They all score 50.

### Criterion B

**Sakai:** a) In its current form Sakai is primarily an e-Collaboration framework, in that there are facilities for online collaboration pre-bundled. There is also currently a software patch that allows Sakai to pick up X.509 certificates from MyProxy servers for Grid computing use. Sakai thus scores 60 for the e-Science sub criterion.

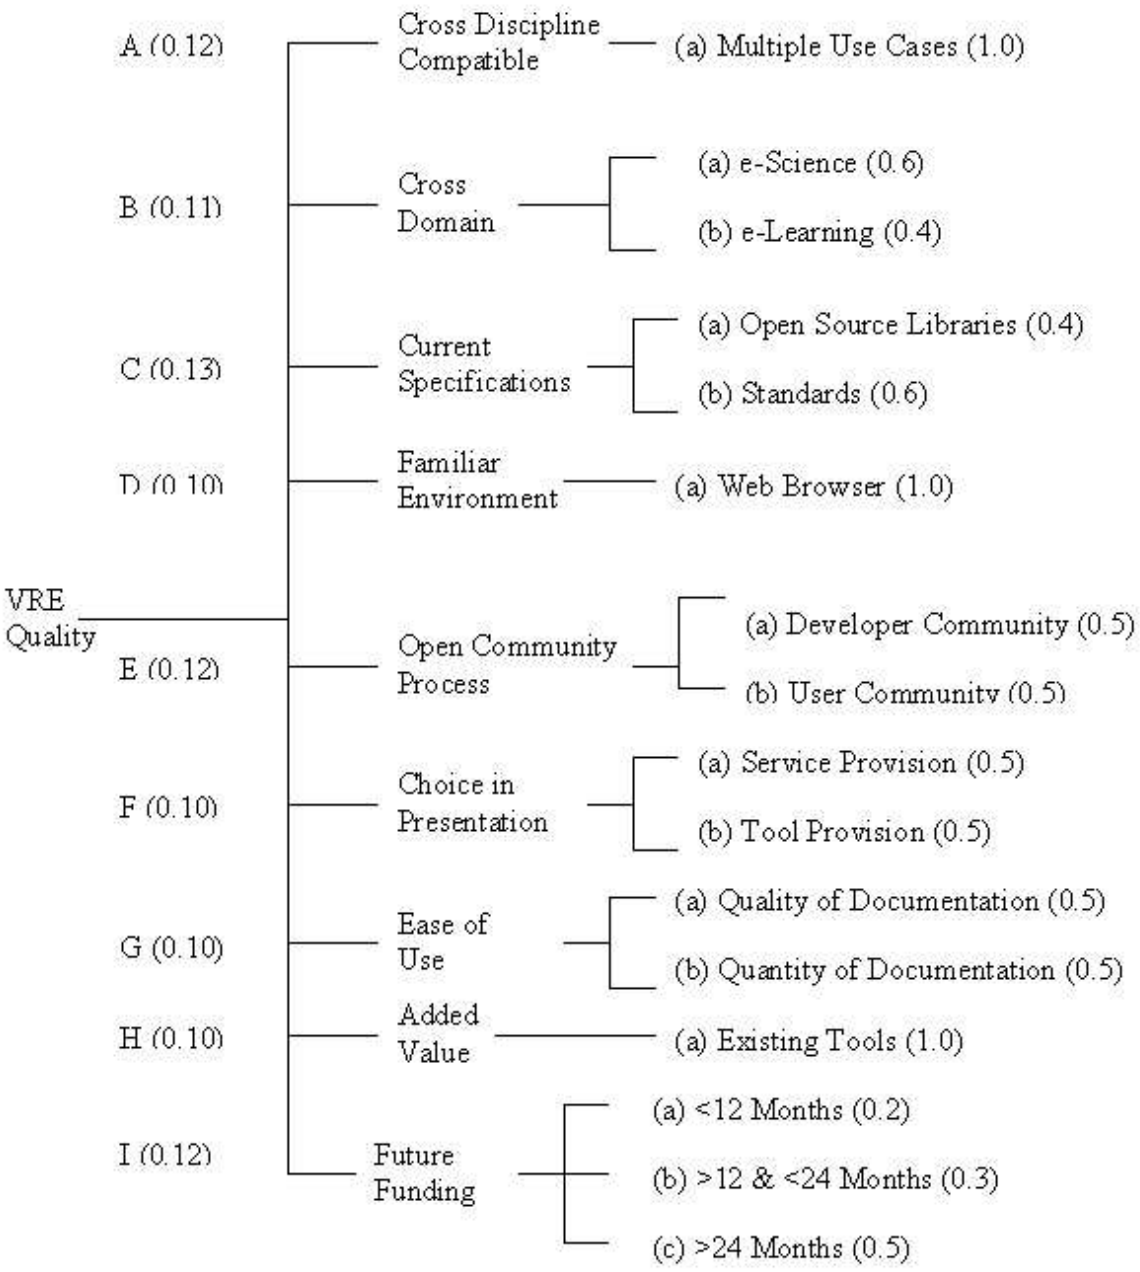


Figure 1: Value tree for comparing open source, platform independent VREs

b) Sakai can also be used for e-Learning, as CHEF is at U. Michigan. It doesn't yet offer any specialized tools for things like assessment and IMS package delivery, but it does offer a working paradigm that lends itself well to representing courses and the students on those courses. Tutorial sessions can be held online, using the bundled collaboration tools, and course materials can be presented in the form of the multiple web pages that you can assign to each course worksite. Sakai scores 50 in the e-Learning category.

**CHEF/ OGCE:** a) The OGCE framework, based on a populated version of CHEF, is the clear winner in the e-Science category. OGCE makes use of CHEFS e-collaboration tools and adds grid computing tool components. NEESGrid in the USA use it in their earthquake simulation experiments. OGCE scores 70 in the e-Science category.

b) CHEF/ OGCE scores slightly less then Sakai in the e-Learning category as it offers similar facilities. See the Sakai section for more details. CHEF/ OGCE scores 45 here.

**Bodington:** a) Bodington is primarily of use as a learning environment. It has strong content management facilities and has tools for generating and displaying assessments complying with IMS QTI (Question and Test Interchange) format. It is, however, not of as great use to the e-Science community, who need collaboration tools and user grouping and administration features more than content navigation. Bodington scores 30 here.

b) This is where Bodington is more fitting. Bodington is a content manager by design, it uses a library metaphor to arrange content into floors and rooms; this is obviously a successful approach as Bodington has been in production at Leeds for around 4 years. A drawback of the Bodington approach is in the centralization of content as opposed to a strategy of organizing metadata on remote content. Bodington scores 70 here.

**GridSphere:** a) GridSphere is a portal framework in the sense that it is an unspecialized container for components that, when added will increase its utility within a specified domain. The main take-up of GridSphere has been within the hard science community and several projects have successfully added portlet components to GridSphere and put it into production. At the GridSphere website you can download a collection of Grid portlets for accessing Globus functionality. GridSphere scores 70 here due to its good Globus Grid tool support.

b) GridSphere scores poorly in the e-Learning sub criterion. There are no tools bundled, or in the pipeline, that are applicable to this domain. GridSphere is a framework designed for specialization however, so the potential is there to add e-portfolio, course management and assessment tools. GridSphere scores 20.

### Criterion C

**Sakai:** a) Sakai makes good use of open source libraries from other projects. It uses the Spring framework, Hibernate, Java Server Faces, Pluto and Velocity amongst many others. It is also re-using lots of the CHEF tool code, which has been tested in the live environment of U. Michigan. For these reasons Sakai scores well and gets 70.

b) Sakai uses the OKI OSIDs for component messaging internally, although these interfaces are not exposed to external tools. There is no support for WSRP or JSR 168 tools presently, JSR 168 support has been in the pipeline for some time and WSRP is being worked into the architecture

design at the time of writing. In reality, the work on formalizing the software interfaces used in Sakai is still ongoing

**CHEF/ OGCE:** a) CHEF uses open source libraries from other projects. It uses Struts, Velocity. JUnit and James, amongst others. CHEF scores 60. On b) CHEF scores 20.

**Bodington:** a) Bodington is a Web application written using servlets for its dynamic aspects. Bodington does not leverage any Web application development frameworks like Java Server Faces, Struts or Tapestry, primarily for the reason that the architecture was effectively formulated in 1998, before the advent of these frameworks. Bodington also doesn't make use of any software libraries that aren't included in the Tomcat installation. As this metric depends on the utilization of cross-project libraries, Bodington scores 30.

b) Bodington's modularity comes at the Java Servlet level. If you want a new tool, you write a servlet. Data repository (DR) services are provided, but are not, to the authors' knowledge, currently based on ratified standards. As with all the other frameworks discussed, you cannot aggregate either JSR 168 or WSRP compliant portlets, so tool creation has to be done by proprietary means. Bodington scores 20 here.

**GridSphere:** a) Gridsphere scores 80. On b) GridSphere scores 50.

### Criterion D

All the frameworks considered score 100 here as they can be displayed in any CSS and JavaScript enabled Web browser.

### Criterion E

**Sakai:** The Sakai project has a formalized community process in place called the SEPP (Sakai Educational Partners Program). The need for the SEPP was actually expressed by the main Sakai funders, the Mellon Foundation, at the Sakai project's inception, and as such is a pre-requisite for funding to continue. To join the SEPP, institutions have to pay a fee of $10K per year for three years. Now that version 1 has been publicly released, there are several public mailing lists for developers and users outside the SEPP. The mailing lists have been dedicated slight more to developers, so Sakai scores 70 and 60 respectively.

**CHEF/ OGCE:** The OGCE project has a good selection of majordomo mailing lists, and a large contact list on its Web site. OGCE scores 60 in each category.

**Bodington:** Bodington is a SourceForge hosted project and has a mailing list dedicated to developers. Traffic is relatively high and this reflects the list's utility. It would be useful if the lists were broken down into more specialized subject areas as there is a lot of material on them. Bodington scores 60 and 30 respectively.

**GridSphere:** GridSphere has both a user's and developers mailing list and so scores 60 in each criterion.

### Criterion F

All the platform tested are presentation agnostic, so they all score 50.

### Criterion G

**Sakai:** Sakai's documentation is patchy and still under development. Seeing as the architecture is still in flux, this is understandable. The state of the documentation is made up for to a degree by the quality of Sakai's community efforts. Sakai scores 20 in each sub criterion.

**CHEF/ OGCE:** OGCE's documentation is better, and more copious than Sakai's, as befits a longer running project. The fact that CHEF long since standardised on its architecture obviously helps. CHEF gets 50 in each category.

**GridSphere:** GridSphere's documentation scores 60 for both quality and quantity; there is a collection of documents in the form of "howtos", both on the Web site and in the binary distribution.

**Bodington:** Bodington scores 40 in each sub criterion. The Bodington gatehouse site hosts a collection of tutorials for students wishing to use the system but these are obviously targeted to its use as a VLE not a VRE.

### Criterion H

CHEF/ OGCE is the clear winner here as it comes with all the collaboration tools of Sakai, plus a suite of useful Grid tools. Sakai comes second, as it comes supplied with a set of useful collaboration tools such as chat, threaded discussion and shared file space. GridSphere comes third as it comes with chat, file space and chart creation tools. Bodington comes last in this criterion, due to the fact that the out of the box toolset relevant to a VRE is sparse, as its intended role is that of an e-Learning/ content management environment, not a tool delivery framework.

### Criterion I

**Sakai:** Sakai funding has only just got going, so we have given it 60 for each sub-criterion.

**CHEF/ OGCE:** CHEF is effectively undergoing a transformation into Sakai and it is unclear what the funding situation currently is. CHEF is still in production use at U. Michigan. CHEF gets 50,10,0 in each sub criterion respectively.

**Bodington:** Bodington is funded by individual project grants from funding bodies, it does not have a concrete funding strategy like Sakai. There is however considerable institutional commitment in the UK and this should really be considered in a funding assessment. As far as we are aware there are no Bodington projects funded past two years scores, so it scores 50, 50 and 20 respectively.

**GridSphere:** GridSphere funding is about to end. Furthermore, GridSphere is currently trying to get funding from Deutsche-Grid initiative to continue support for GridSphere, but now it looks like it may be only enough money to support only 1 developer for 3 years in the worst case, so we have only given them 20 in each sub-criterion.

## 2.4 Combining the scores and weights

### Results

We also performed a sensitivity analysis (1) by changing the weights, (2) having someone else decide

| | | | | | | | | | | | | | | | | | Total | Ranks |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| **Criteria** | **A** | **B** | | **C** | | **D** | **E** | | **F** | | **G** | | **H** | **I** | | | | |
| **Criteria weights** | 0.12 | 0.11 | | 0.13 | | 0.1 | 0.12 | | 0.1 | | 0.1 | | 0.1 | 0.12 | | | 1 | |
| **Sub Criteria** | a | a | b | a | b | a | a | b | a | b | a | b | a | a | b | c | | |
| **Sub Weights** | 1 | 0.6 | 0.4 | 0.4 | 0.6 | 1 | 0.5 | 0.5 | 0.5 | 0.5 | 0.5 | 0.5 | 1 | 0.2 | 0.3 | 0.5 | | |
| **Sakai 1** | 50 | 60 | 50 | 70 | 30 | 100 | 70 | 60 | 20 | 30 | 20 | 20 | 60 | 60 | 60 | 60 | 53.64 | 1st |
| **Chef/OGCE** | 50 | 70 | 45 | 60 | 20 | 100 | 60 | 60 | 20 | 30 | 50 | 50 | 70 | 50 | 10 | 0 | 50.54 | 2nd |
| **Bodington** | 50 | 30 | 70 | 30 | 20 | 100 | 60 | 30 | 20 | 30 | 40 | 40 | 30 | 50 | 50 | 20 | 43.28 | 4th |
| **GridSphere** | 50 | 70 | 20 | 80 | 50 | 100 | 40 | 50 | 20 | 30 | 60 | 60 | 40 | 20 | 20 | 20 | 49.86 | 3rd |

Figure 2: Combining Scores and Weights

on their own scores, (3) by dropping the least important criteria (D,F,G and H). In each situation we obtain very similar ranks, Sakai always comes out on top.

## 3 Acknowledgements

We acknowledge the following people who have freely provided their comments and input to the work reported here.

Alison Allden, Chris Awre, Mark Baker, David Boyd, David de Roure, Geoffrey Fox, Dennis Gannon, Glen Golden, Liz Lyon, Mark Norton, Jason Novotny, Marlon Pierce, Andy Powell, Alan Robiette, Charles Severance, Mary Thomas, Brad Wheeler, Scott Wilson,

We especially thank Mary Thomas, Dennis Gannon, Jason Novotny, Charles Severance, Glen Golden, Brad Wheeler and Jim Farmer for their continual guidance and help.

## References

[1] R. Crouchley, R.J. Allan et al. *A Service-Oriented Architecture for VRE/ VLE building on the Sakai Framework* Report to JCSR, JCSR(04)14a (2004)

[2] Mark Baker, Hong Ong, Rob Allan and Xiao Dong Wang *Advanced Portal Services* Proc. AHM'2004 (Nottingham, September 2004) `http://www.allhands.org.uk/proceedings/papers/223.pdf`

[3] Rob Allan, Rob Crouchley, Adrian Fish and Dharmesh Chohan *Portal Services for Awareness and Training* Proc. AHM'2004 (Nottingham, September 2004) `http://www.allhands.org.uk/proceedings/papers/234.pdf`

[4] W. Edwards, and J.R. Newman *Multiattribute Evaluation* (Sage, 1982)

use of information contained in any of their reports or in any communication about their tests or investigations.